Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey

By

Daron Acemoglu
Massachusetts Institute of Technology

Gary Anderson
National Center for Science and Engineering Statistics

David Beede
U.S. Census Bureau

Catherine Buffington
U.S. Census Bureau

Eric Childress
George Mason University

Emin Dinlersoz
U.S. Census Bureau

Lucia Foster
U.S. Census Bureau

Nathan Goldschlag
U.S. Census Bureau

John Haltiwanger
University of Maryland

Zachary Kroff
U.S. Census Bureau

Pascual Restrepo
Boston University

Nikolas Zolas
U.S. Census Bureau

CES 22-12 April 2022

Abstract

This paper provides a comprehensive description of the adoption of automation technologies by US firms across all economic sectors by leveraging a new module introduced in the Census Bureau's 2019 Annual Business Survey. The module collects data from over 300,000 firms on the use of five advanced technologies: AI, robotics, dedicated equipment, specialized software, and cloud computing. We document that the adoption of these technologies remains low (especially for AI and robotics), varies substantially across industries, and concentrates on large and younger firms. However, because larger firms are much more likely to adopt them, 12-64% of US workers and 22-72% of manufacturing workers are exposed to these technologies. Firms report a variety of motivations for adoption, including automating tasks previously performed by labor. Consistent with the use of these technologies for automation, adopters have higher labor productivity and wages and lower labor shares. In particular, the use of these technologies is associated with a 15% increase in labor productivity, which accounts for 20–30% of the higher labor productivity achieved by the largest firms in an industry. Adopters report that these technologies raised skill requirements and led to greater demand for skilled labor, but brought limited or ambiguous effects to their employment levels.

* Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the US Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this data product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. DRB Approval Numbers: CBDRB-FY21-058, CBDRD-FY21-316, CBDRB-FY22-057. We thank Laurence Ales, Chiara Criscuolo, Eric Donald, Christina Patterson, and participants in the 2021 AEAStat session, 2021 Meetings of the Society for Economic Dynamics, and NBER CRIW conference for comments and suggestions. Acemoglu gratefully acknowledges financial support from the National Science Foundation, the Hewlett Foundation, Schmidt Sciences, and the Smith Richardson Foundation. Restrepo thanks the National Science Foundation for its support under award No. 2049427.

1 Introduction

Advanced technologies, including robotics, artificial intelligence (AI), and various software products, are thought to be spreading rapidly in industrialized economies.[1] These technologies are often argued to increase productivity, automate tasks performed by labor, and raise the demand for skills, contributing to rising inequality and declining labor shares. There is little systematic evidence, however, on how widespread these technologies are, which firms are adopting them, and how they affect firm and worker outcomes, especially in the US.

This paper leverages a new module introduced in the 2019 Annual Business Survey (ABS) conducted by the US Census Bureau in partnership with the National Center for Science and Engineering Statistics (NCSES) to shed light on these questions. The module focuses on the use of five advanced technologies: artificial intelligence (AI), robotics, dedicated equipment, specialized software, and cloud computing. The 2019 ABS module sampled over 300,000 employer businesses and collected information on firms' adoption and use of these technologies for the 2016–2018 reference period. The module also asked questions on firms' motivation for adoption, allowing us to measure the extent to which these technologies are being used for automating labor-intensive tasks, and on firms' assessments of the impact of technology on the size and skills of their workforce.

Using the ABS data, this paper provides a comprehensive description of the adoption of advanced technologies by US firms across all economic sectors, documenting for the first time the extent to which these technologies are being used for automation and how they affect firms' production processes and demand for skills.

We first show that, even though a minority of firms use these technologies, these firms account for a sizable share of employment. In particular, only 3.2% of firms currently use AI as part of their processes and methods and 2% use robotics. There is wider diffusion of the remaining technologies, with 19.6% of firms using dedicated equipment, 40.2% using specialized software, and 34% using cloud computing. Still, half of US firms used none of these technologies by 2018. However, because adoption concentrates among the largest firms in the US economy, a sizable fraction of the workforce is exposed to these technologies: 12.6% of US workers are employed in firms using AI between 2016-2018 (even though these are only 3.2% of US firms), and the analogous shares are 15.7% for robotics, 36.4% for dedicated equipment, 64.4% for specialized software, and 61.8% for cloud computing. These high exposures indicate that advanced technologies have been an important force affecting US labor and product markets, despite their limited diffusion across firms, especially among the smallest ones.

The low adoption rates across firms appear related to the lack of applicability and high adoption costs of these technologies. In particular, to explore why the majority of firms have not invested in advanced technologies, the ABS also asked firms about factors limiting their adoption. A large fraction of non-adopters recognizes that these technologies are not applicable to their business models, which points to significant heterogeneity in the extent to which firms can benefit from these

[1] See Brynjolfsson and McAfee (2014), Ford (2015), Susskind and Susskind (2015) and Schwab (2017).

technologies. Conditional on the technology being applicable, both adopters and non-adopters cite the high cost of adoption as the main factor limiting wider adoption, which points to sizable fixed costs involved in deploying and integrating these technologies, and explain the limited adoption by small firms.

The ABS module allows us to document for the first time the extent to which US firms are using advanced technologies for automation and the importance of this use for the workforce. Firms report a number of motivations for their investments in these advanced technologies, including most commonly improving process quality, upgrading existing processes, and automating tasks performed by labor. The use of AI and robotics is closely related to automation, with 55% and 65% of users (in an employment-weighted sense) adopting these technologies for automation, respectively. Dedicated equipment and specialized software, on the other hand, have more diverse uses, with 30–35% of users adopting these technologies for automation. Other uses of these technologies, such as expanding their product offerings or meeting industry standards are less common.

As a whole, the fraction of the US workforce exposed to automation-related uses of advanced technologies is sizable, with 7–20% of US workers employed at firms using advanced technologies for automation. Worker exposure to automation is particularly high in (though not exclusive to) manufacturing, where 17–37% of workers are employed in firms using these technologies for automation.[2] Although AI and robotics stand out as the two technologies whose use is more closely related to automation, most of workers' exposure to automation comes from dedicated equipment and specialized software due to the wider diffusion of these technologies.

After documenting these aggregate facts, we turn to exploring differences in adoption rates across industries and across firms within an industry. Adoption rates vary substantially across industries, which is in line with firms' reports that these technologies have highly specific capabilities. Detailed industry differences account for 10-30% of the (employment-weighted) variation in adoption rates across businesses. While adoption rates are higher in manufacturing, we also document sizable heterogeneity among detailed manufacturing industries. For example, in the case of robotics, some industries like forging (NAICS 3221) and motor vehicle manufacturing (NAICS 3363) have some of the highest adoption rates in this sector.

The ABS also points to vast differences in adoption across firms *within* the same detailed industry. Adoption concentrates in larger and younger firms, presumably reflecting the large fixed costs and organizational barriers involved in adopting these technologies. Consistent with the importance of automation as a major application of these technologies, adopters have higher labor productivity and lower labor shares than other firms in their same industry, size class, and cohort.[3]

[2] When interpreting these measures, one should keep in mind that not all workers currently employed at firms using these technologies are subject to the effects of automation. Thus, our numbers do not imply that 7–20% of US workers are or will be at risk of having their jobs automated. Nevertheless, these high shares suggest that changes in the demand for skills at automating firms will likely have sizable implications for certain groups of workers throughout the economy (see Acemoglu and Restrepo, 2021b).

[3] This is in line with recent papers that emphasize the role of automation as a key driver behind the labor share decline in some sectors, most notably in manufacturing (see Acemoglu, Lelarge and Restrepo, 2020; Acemoglu and Restrepo, 2021b; Dauth et al., 2021; Cheng et al., 2021; Kogan et al., 2021).

Also consistent with the incentives for automation increasing when wages are higher, we find that higher-wage firms are more likely to adopt these technologies.[4]

The fact that the adoption of advanced technologies is associated with higher labor productivity and that the adoption of these technologies concentrates in large firms suggests that these technologies contributed to the *superstar firm* phenomenon—the rise of large firms with high labor productivity and low labor shares—documented by Autor et al. (2020) and others. Using the ABS data we show that the use of advanced technologies is associated with a 15% increase in labor productivity (and a 26% higher labor productivity for firms using all five technologies surveyed in the ABS). Coupled with the fact that adoption concentrates at large firms, these estimates imply that, from a pure descriptive viewpoint, the adoption of advanced technologies accounts for 20–30% of the labor productivity differences between small and the superstar firms in each industry.[5]

Finally, we explore firms' self-assessment on the implications of advanced technologies for their demand for labor and skills. Most firms report that advanced technology adoption did not change their overall employment levels, and among firms that report a change the findings are split, which point to limited and ambiguous effects of advanced technologies on firm employment levels. Instead, a significant share of firms (between 30–40% depending on the technology surveyed) assess that advanced technologies increased their skill demands, while almost no firms report a reduction in their demand for skills. These self-reports are consistent with theories emphasizing that advanced technologies increase the demand for skills. However, the evidence does not support the notion that the higher labor productivity brought by these technologies translates into a uniform increase in employment levels.[6]

Our paper contributes to a growing literature on measuring the adoption of advanced technologies across firms and industries and understanding their implications for firms and workers. Our first contribution is on data collection and the measurement of technology adoption. In this we are building upon and expanding the work in the 2018 ABS, summarized in Zolas et al. (2020a). Due to data limitations, other prior research on the effects of modern automation technologies on firms and workers has relied on indirect proxies of technology, or on datasets with limited and coarse coverage. Earlier work focused on industry-level robot adoption measures, from the In-

[4]See Acemoglu and Restrepo (2021a) and Dechezlepretre et al. (2021) for evidence on how high wages drive the adoption of automation technologies across countries and firms, respectively.

A different interpretation is that automation changes the composition of the workforce and pushes firms to increase their demand for high-wage high-skilled workers.

[5]This is consistent with the results of Acemoglu, Lelarge and Restrepo (2020) from French manufacturing, which indicate that a sizable portion of the covariance between labor share changes and size is related to robotics investments. This also aligns with the conclusions in Hubmer and Restrepo (2021), who use a calibrated model of firm dynamics with differences a fixed cost of technology adoption to show that the uneven adoption of automation technologies by large firms contributed to the rise of superstar firms.

[6]We also caution that, as documented in Koch, Manuylov and Smolka (2021) and Acemoglu, Lelarge and Restrepo (2020), the expansion of firms adopting automation technologies might come at the expense of competing firms that do not automate. Hence the increases in employment and the lack of negative employment effects reported by adopters is consistent with potentially negative industry- or market-level effects as found in a number of studies, such as Acemoglu and Restrepo (2020), Dauth et al. (2021), Acemoglu, Lelarge and Restrepo (2020), or no positive aggregate effects, as in Graetz and Michaels (2018).

ternational Federation of Robotics.[7] A more recent series of papers uses data on robot imports and in some cases detailed surveys of manufacturing firms in order to explore firm-level outcomes for robot adopters in manufacturing.[8] Our paper extends these efforts by collecting and utilizing comprehensive data for the entire economy, not just manufacturing, and including AI, dedicated equipment, specialized software, and cloud computing as well as robotics. In doing so, we confirm and extend four findings that many of these papers have documented for robots and extend it to the other technologies in the ABS: (i) larger firms are much more likely to adopt these advanced technologies; (ii) adoption is associated with higher labor productivity; (iii) adoption is associated with lower labor shares; and (iv) advanced technology adoption is associated with an increasing demand for skills (for example, in the form of reductions in the share of production workers for robotics).

There is also a nascent literature using various proxies for firm-level adoption of AI and new technologies obtained from the text in job postings, conference calls, and patent data. For example, Alekseeva et al. (2021); Babina et al. (2021) and Acemoglu et al. (2020) use data from online vacancies, specifically from postings including AI-related skills, to estimate establishment-level AI activity. Bloom et al. (2021) combine the job-postings data with the text data from patents and conference calls to measure the diffusion of AI and other novel technologies across US firms. Mann and Püttmann (2019), Dechezleprêtre et al. (2021), and Martinez and Moen-Vorum (2021) use the text of patents to identify and measure the diffusion of automation technologies. The ABS module contributes to these efforts as it is more representative and offers a more direct measurement of technology use at the firm level. The ABS also distinguishes between users and providers of technology—a distinction that approaches based on patent data and job postings miss.

The ABS technology module also complements new surveys that have been recently used to measure technology adoption in other countries. For example, there are similar surveys measuring firm's investments in advanced technologies and automation for Germany (see Genz et al., 2021), Italy (see Calvino et al., 2022), South Korea (see Cho et al., 2021), the Netherlands (see Bessen et al., 2019) and manufacturing in Spain (see Koch, Manuylov and Smolka, 2021).

Finally, we also expand on works using previous Census products devoted to measuring technology adoption in specific sectors. These include the earlier Survey of Manufacturing Technology, which was conducted in 1988, 1991, and 1993 and collected data on robots, automated storage and retrieval systems, automated guided vehicle systems, and automated testing equipment for a subset of manufacturing industries (Doms, Dunne and Troske, 1997; Dinlersoz and Wolf, 2018); computer use in manufacturing using the Annual Survey of Manufacturers (Dunne et al., 2004); telecommunication technologies using the Information and Communication and Technology Survey (ICTS) (Eckert, Ganapati and Walsh, 2020); e-business practices using the Computer Network

[7] See, for example, Graetz and Michaels (2018) and Acemoglu and Restrepo (2020).

[8] See for example Humlum (2020); Bonfiglioli et al. (2020); Rodrigo (2021); Dixon, Hong and Wu (2020); de Souza and Sollaci (2020) for papers using data on robot imports to measure firm-level adoption of robotics in various countries. Acemoglu, Lelarge and Restrepo (2020) combine imports data and detailed accounting information to construct a somewhat more complete picture of robot adoption in French manufacturing. Aghion et al. (2020) also complement these data with accounting information on the use of motor equipment.

Use Supplement (CNUS) (McElheran, 2015); customer–employee substitution at gas stations using the Census of Retail Trade Basker, Foster and Klimek (2017); and manufacturers use of data and predictive analytics using the Management and Organizational Practices Survey (MOPS) which collects detailed information on firm organization and management practices (Brynjolfsson, Jin and McElheran, 2021).

The rest of the paper is organized as follows. Section 2 describes the ABS module. Section 3 outlines the conceptual framework that motivates our empirical work and interpretation. Sections 4, 5, 6, and 7 provide our analysis of the ABS data. Section 8 concludes, while the Appendix includes additional details on the design of the ABS module and empirical results.

2 The Technology Module from the 2019 ABS

The 2019 technology module is the second module on technology collected as part of the ABS. The first technology module featured in the 2018 ABS focused on questions regarding firms' digitization of information along with the adoption of some specific business technologies (see Zolas et al., 2020b, for an analysis of the data collected by the 2018 ABS technology module). The new 2019 ABS module features questions relating to the adoption of five advanced technologies that are relevant for automation. In addition, the 2019 ABS module collects data on firms' motivations for adoption, asking firms to report whether they are using these advanced technologies for automation, and their assessment of the effect of these technologies on the size and composition of their workforce. Finally, the module asks firms about the bottlenecks limiting their adoption of advanced technologies.[9]

The five technologies surveyed in the 2019 ABS were:

- **Artificial Intelligence:** Artificial intelligence is a branch of computer science and engineering devoted to making machines intelligent. Intelligence is that quality that enables an entity to perceive, analyze, determine response and act appropriately in its environment.
- **Robotics:** Robotic equipment (or robots) are automatically controlled, reprogrammable, and multipurpose machines used in automated operations in industrial and service environments.
- **Specialized Software (excluding Artificial Intelligence):** Specialized software is software dedicated to performing a particular business function.
- **Dedicated Equipment (excluding Robotics):** Specialized equipment is equipment capable of automatically carrying out pre-specified task(s).
- **Cloud-based Computing Systems and Applications:** Cloud systems and applications are computing resources available on-demand via the internet.

[9] Appendix A provides an overview of the development of the 2019 ABS module.

According to these definitions, AI algorithms powering a chatbot would count as **artifical intelligence**. Industrial robots used in manufacturing would be considered **robotics**. A software system for document discovery or handling appointments would be **specialized software**. And an automatic retrieval system for warehouses would be **dedicated equipment**. On the other hand, **cloud computing** is typically used together with the other technologies to satisfy their demand for computing power. For example, a firm using a software system for trading would host the algorithms and data on the cloud.

The module starts by asking firms about their use of these technologies during the reference period of 2016–2018.[10] For each technology, a firm may respond that it did not use the technology; tested the technology, but did not use it; used the technology with a specified degree of intensity (low use, moderate use, or high use); or doesn't know whether the technology was used during the three years 2016–2018. Conditional on responding with some degree of technology use (low, moderate, or high), firms are then asked about the motivations for adopting or using the technology. Respondents may choose from the following list of motivations: "automate tasks performed by labor", "upgrade outdated processes or methods", "improve quality or reliability of processes or methods", "expand the range of goods or services", "adopt standards and accreditation", and "some other reason". In addition, the module asks firms whether they are using the technologies as part of their production processes and methods, or whether they are providers of these technologies (or goods and services that embed these technologies).

The module then dives into the workforce effects of technology. First, respondents reporting some degree of technology use are asked about the effects of the technology on overall employment, overall skill level, and STEM skills of their workers. Firms may respond with "increased", "decreased", or "did not change" ("not applicable" is also an option for STEM skills in the case that the firm did not employ any workers with STEM skills). Next, firms are asked about how technology use affected four types of workers—production, non-production, supervisory, and non-supervisory. Firms can respond that technology use either increased, decreased, did not change the number of each of these types of workers (again, with the additional option "not applicable"). These questions, while qualitative in nature, provide a broad assessment of various effects of the technology on a firm's workforce size and composition.

Finally, all firms—regardless of reporting technology use or not—were asked to assess factors adversely affecting adoption and utilization of each technology. The factors included represent a wide variety of considerations and concerns, including applicability/feasibility ("technology not applicable to this business"), the cost and maturity of technology ("this technology was too expensive", "this technology was not mature"), inputs needed to utilize the technology ("lacked access to required data","required data not reliable", "lacked access to capital", "lacked access to required human capital and talent"), regulatory environment ("laws and regulations"), and security considerations ("concerns regarding safety and security—physical and cyber"). A response option of

[10] The exact wording and organization of the questions is available in the official survey instrument for 2019 ABS on the ABS website: https://www.census.gov/programs-surveys/abs/technical-documentation/surveys-instructions.2019.html.

unhindered adoption and utilization was also included ("no factors adversely affected the adoption of this technology").

The set of firms sampled in the technology module was dictated by the general sampling scheme for the 2019 ABS, a primary goal of which is to provide tabulations of collected data by various ownership characteristics.[11] The ABS sampling universe was created using Census Bureau's Business Register administrative data from 2018, which provides the information on industry classification, receipts, payroll and employment for the construction of ABS universe. The ABS universe was stratified by state, frame, and industry, where *frame* refers to categories of ownership characteristics for businesses. The Census Bureau used several sources of information to estimate the probability that a business is minority or women-owned. These probabilities were then used to place each firm in the ABS universe to one of nine frames that span key race and ethnicity categories, plus gender and public ownership status. Large companies were selected with certainty based on volume of sales, payroll, or number of paid employees.[12] The remaining universe was subjected to stratified systematic random sampling.

The 2019 ABS data were collected from June through December 2019. The response rate for the portion of the survey used in this paper was 68.7%, which is in line with the 2018 ABS technology module response rate. The subset of sampled firms that responded to the technology module does not constitute, even after weighting by ABS sampling weights, a nationally representative distribution of firms, when compared to the full set of employer businesses in the Longitudinal Business Database (LBD). The LBD contains the universe of non-farm employer business establishments in the US (see Chow et al., 2021, for details). We therefore construct firm weights based on the 2018 LBD to make the sample representative of the universe of employer businesses. These weights are calculated using a methodology similar to the one used in Zolas et al. (2020b). Specifically, we first stratify firms in the 2018 LBD and 2019 ABS by the same size, age, and industry categories (12 size categories, 12 age groups and 19 two-digit NAICS sectors). Each firm in a stratum in the ABS is then assigned the same weight calculated by dividing the firm count in the corresponding 2018 LBD stratum by the firm count in the 2019 ABS stratum.

An additional challenge is item non-response. Many firms did not respond to portions of the module or specific questions. However, the item response rate of the technology module does not differ substantially from that of the remaining part of the ABS. Another concern, as in the 2018 ABS technology module, is that some firms (especially large and old firms) responding not knowing whether they used some of the technologies in the survey. We remove firms that respond "Don't Know"from our analysis of that technology.

[11]For details on the sampling methodology, see https://www.census.gov/programs-surveys/abs/technical-documentation/methodology.2019.html.

[12]More specifically, certainty cases satisfy the following criteria: firms with more than 500 employees; firms responding to the 2016 Business R&D and Innovation Survey for Microbusinesses (BRDI-M) survey with R&D costs of $1 million or higher; and firms larger than stratum-specific payroll and receipt cut-off. The certainty cutoffs vary by stratum, depending on the number and the size distribution of firms in the stratum.

3 Conceptual Framework

This section sketches a partial-equilibrium version of the model in Acemoglu and Restrepo (2021b), expanded to include firm competition within an industry, in order to frame our interpretation of the results from the ABS technology module. In our framework, firms complete tasks to produce output, and their key decision is the assignment of these tasks across workers with different skills and specialized capital equipment or algorithms (e.g., AI, robotics, specialized software, etc.). Automation is the use of specialized capital in order to perform tasks previously assigned to labor. The results presented in this section follow from those in Acemoglu and Restrepo (2021b) and their proofs are omitted.

3.1 Production, Tasks and Demand

We consider a partial equilibrium model of a single industry i. To save on notation, we omit industry subscripts, with the understanding that all objects and technologies might vary by industry. Firms are indexed by f and engage in monopolistic competition, facing a demand curve for their products given by

$$y_f = y \cdot \left(\frac{p_f}{p}\right)^{-\sigma}, \quad \text{with } \sigma > 1 \tag{1}$$

and charge a constant markup of $\mu = \sigma/(\sigma - 1)$. Here, y is industry output and p the industry price index.

Output is produced by completing a mass M of tasks indexed by x and belonging to some set $\mathcal{T}$. The production function for firm f is

$$y_f = z_f \cdot \left(\frac{1}{M}\int_{\mathcal{T}} (M \cdot y_f(x))^{\frac{\lambda-1}{\lambda}} \cdot dx\right)^{\frac{\lambda}{\lambda-1}},$$

where z_f denotes the (factor-neutral) productivity of the firm, $y_f(x)$ denotes the quantity of task x completed, and $\lambda > 0$ is the elasticity of substitution between tasks.

All firms in the industry complete the same set of tasks, $\mathcal{T}$, but differ in their productivity z_f, the factor prices they face, and how they assign tasks to different factors. Specifically, task x can be performed using workers from different skill groups, indexed by g,

$$y_f(x) = \sum_g A_g \cdot \psi_g(x) \cdot \ell_{g,f}(x).$$

Here, $\ell_{g,f}(x)$ is the quantity of labor of type g employed by the firm at task x, A_g denotes the productivity level of these workers across all tasks, and $\psi_g(x)$ denotes their productivity at task x. Task-specific productivities $\psi_g(x)$ capture the comparative advantages of groups of workers across tasks (e.g., less educated workers might have a comparative advantage in manual tasks).

Firm f can pay a fixed cost $\kappa_f(x)$ to adopt and integrate the technology required to automate the production of task x. Depending on the task, this may involve the use of AI, robotics, dedicated equipment or specialized software to complete the task. We denote the specialized capital used for this purpose by $k_f(x)$, which yields the automated production of task x for firms that pay the fixed cost as

$$y_f(x) = A_k \cdot \psi_k(x) \cdot k_f(x),$$

where A_k gives capital productivity across all tasks and $\psi_k(x)$ is the productivity of specialized capital at task x. The fact that automation takes place at the task level, using task-specific specialized equipment or algorithms, aligns with the way in which industrial robots, narrow AI algorithms, dedicated equipment and specialized software are used in practice. This also underscores that automation technologies are not "general-purpose", but rather highly application and task-specific.

We close the model by assuming that firms pay exogenous wages given by $w_{f,g} = \tau_f \cdot w_g$, where w_g is the common component of the wage for workers in skill group g, faced by all firms in the industry, while τ_f is a firm-specific component, reflecting differences in the labor supplied faced by firms or the way they share rents with workers. On the other hand, we take the user cost of capital, w_k, to be common across firms and tasks.

3.2 Costs of Automation and Differences in Technology Across Firms

In the benchmark case in which wages are identical across firms and there are no fixed costs of automation, all firms would make the same cost-minimizing automation decisions. Differences in the factor-neutral productivity term z_f do not impact automation decisions and simply translate into differences in firm scale. In this benchmark, despite having different scales, all firms in an industry would employ the same bundle of workers, machines, and software to produce, and would have the exact same labor productivity and labor share.

As described in the Introduction, however, there are sizable differences in adoption rates across firms, both between and within industries. We view these differences as being due to three factors:

- First, the nature of tasks required in an industry and firm determines the productivity and applicability of different types of automation technologies. For example, industrial robots are not useful in most non-manufacturing industries, and within manufacturing, they are most suitable for various manual tasks involved in heavy industry, such as welding, painting, sorting and assembly, while certain fine-motor tasks, such as stitching of shoes, are harder for robots.[13] Yet other manual tasks, involved in spinning, weaving and stitching in textile industries can be automated using dedicated equipment. Likewise, a range of white-collar tasks in services can be automated using specialized software and increasingly AI.

- Second, even firms that compete in the same industry may face different wages, contributing to differences in adoption decisions.

[13]In line with this argument, data from the 2018 Annual Survey of Manufactures (ASM) published by the US Census show that about 1% of establishments in apparel manufacturing use robots while 35% of establishments in motor vehicle manufacturing do.

- Third, and most importantly, the fixed costs of adopting and integrating automation technologies will preclude some firms from using them. These fixed costs depend on industry (e.g., because it determines the engineering complexity of the tasks to automate), firm age (e.g., because they might face less organizational barriers), and other firm-level characteristics (e.g., how digitally savvy or informed the management may be and whether the firm needs to customize its products). For this reason, younger firms might have lower fixed costs, while the same fixed cost will make automation technologies less profitable for smaller firms in industries with limited applications for advanced technologies. Integration costs can be sizable. For example, in manufacturing, integration costs associated with the use of industrial robots can add up to four times the cost of the actual equipment (see, for example, Leigh and Kraft, 2018).

3.3 Automation, Factor Shares, and labor Productivity

Let $\mathcal{T}_{f,k}$ denote the set of tasks that firm f has automated, while $\mathcal{T}_g$ denotes the set of tasks for which production by worker of type g would minimize costs in the absence of automation. Following Acemoglu and Restrepo (2021b), define the *task share* of workers of type g and capital at firm f as

$$\Gamma_{f,g} = \frac{1}{M} \int_{\mathcal{T}_g - \mathcal{T}_{f,k}} \psi_g(x)^{\lambda-1} \cdot dx, \qquad \Gamma_{f,k} = \frac{1}{M} \int_{\mathcal{T}_{f,k}} \psi_k(x)^{\lambda-1} \cdot dx.$$

The unit cost of production for firm f is then

$$c_f = \frac{1}{z_f} \cdot \left(\sum_g \Gamma_{f,g} \cdot w_{f,g}^{1-\lambda} + \Gamma_{f,k} \cdot w_k^{1-\lambda} \right)^{\frac{1}{1-\lambda}}.$$

Although this unit cost resembles the standard constant elasticity of substitution price index (which would result when firms have CES production functions), the shares are now endogenous and depend on $\mathcal{T}_{f,k}$—the set of tasks that the firm has automated. Input shares are also related to $\mathcal{T}_{f,k}$. As a result, the share of labor in costs, which is proportional to the share of labor in value added, is

$$s_{\ell,f} = \frac{\sum_g \Gamma_{f,g} \cdot w_g^{1-\lambda}}{\sum_g \Gamma_{f,g} \cdot w_g^{1-\lambda} + \Gamma_{f,k} \cdot w_k^{1-\lambda}}.$$

Likewise, the share of labor of type g in the wage bill can be computed as

$$s_{g,f} = \frac{\Gamma_{f,g} \cdot w_g^{1-\lambda}}{\sum_{g'} \Gamma_{f,g'} \cdot w_{g'}^{1-\lambda}},$$

and labor productivity (defined as sales per worker) as:

$$\text{labor productivity}_f = \mu \cdot \bar{w}_f \cdot \frac{\sum_g \Gamma_{f,g} \cdot w_g^{1-\lambda} + \Gamma_{f,k} \cdot w_k^{1-\lambda}}{\sum_g \Gamma_{f,g} \cdot w_g^{1-\lambda}},$$

where $\bar{w}_f$ denotes the average wage paid by the firm. These expressions show that the labor share, the hare of each skill in costs, and labor productivity are all shaped by task shares and thus which tasks are automated.

In summary, labor shares and labor productivity will differ across firms because of variation in task shares. In particular, firms' automation decisions, summarized by the set of tasks they have automated $\{\mathcal{T}_{k,f}\}$, determine their factor shares and labor productivity. In turn, firms will automate different sets of tasks depending on the fixed costs $\kappa_f(x)$, firm-specific wages, and industry differences in the nature of tasks.

3.4 Implications

The framework delivers the following implications:

- The adoption rate of automation technologies could be low because of the high specificity of the tasks that can be effectively automated, high integration costs, or organizational barriers to new technology.

- There will be large differences in adoption across industries driven by the applicability of advanced technologies to production tasks.

- Because of the fixed cost of adoption, large firms are more likely to adopt automation technologies. If fixed costs of adoption, κ_f, are lower for younger firms, then we expect younger firms to also adopt these technologies at a higher rate, conditional on their size.

- All else equal, higher-wage firms—those with larger τ_f—are more likely to adopt automation technologies as well.

- Adopters of automation technologies will have lower labor shares and higher labor productivity.

- Automation, by reducing production costs, always expands firm sales but has an ambiguous effect on firm employment. The overall employment effect depends on whether the productivity effect (the higher sales induced by the cost reduction generated by automation) dominates the displacement effect (the fact that the firm becomes less labor intensive). Independently of which effect dominates, firm employment effects of automation overstate the industry-level implications, since automating firms will expand in part at the expense of their competitors.

- If automated tasks used to be performed by lower-skill groups (which is what we would expect to the extent that routine tasks are more likely to be automated), then advanced technologies will also increase (average) skill requirements directly.

- The implications of different technologies (AI, robotics, etc...) for labor productivity and demand for skills will vary depending on the types of tasks that it automates. Technologies automating tasks performed by workers with lower skill levels will lead to a greater increase in labor productivity and demand for skills.

Our framework also underscores a key distinction between automation and other forms of technological progress. For example, a technology that simply raises productivity in a factor-neutral way—z_f in our model— would also raise firm sales and employment by the same amount, but would have no impact on the labor share, skill requirements or labor productivity (measured as sales per worker).[14] Hence, automation technologies have very distinct effects than factor-neutral technologies on firms' demand for skills, factor shares, and labor productivity. As we will see, the correlations in the data are consistent with the implications we expect from advanced technologies being used for automation, and not simply to increase TFP in a factor neutral way.

4 Aggregate Patterns on the Adoption of Advanced Technologies and the Importance of Automation

This section documents the aggregate patterns on adoption rates by technology as well as the motivations behind the adoption of advanced technologies and the factors limiting a wider diffusion.

4.1 Adoption Rates by Technology

Table 1 reports the share of firms using each of the technologies as part of their processes and methods. Column 1 shows that the share of firms using these advanced technologies is low for AI and robotics (with 3.2% of US firms using AI and 2% using robotics), and higher but still moderate for the remaining technologies (19.6% for dedicated equipment, 40.2% for specialized software, and 34.0% for cloud computing).[15] In total, 47.6% of US firms had adopted at least one of these technologies by 2018, implying that half the firms in the US had not.[16]

Even though a minority of firms use these technologies, these firms are among the largest in the US and account for a sizable share of employment. As a result, the share of workers *exposed* to advanced technologies is significantly higher than the share of firms using these technologies. Column 2 documents this phenomenon. For example, even though only 2% of US firms use robotics,

[14] It might at first be counterintuitive that higher z_f has no effect on labor productivity. To understand this result, first note that higher z_f increases TFP, but labor productivity, defined as (dollar value of) sales divided by labor, is invariant to it. To see this, take the simple example in which there is only one type of labor, no capital, and the firm's production function is simply $z_f \ell_f$, and the firm still faces the demand curve given by (1). In this case, an increase in z_f increases real output per worker, but reduces price by exactly as much, so labor productivity remains constant. As this example clarifies, this holds so long as firms face a demand curve with constant demand elasticity.

[15] The adoption rates of robotics and AI are close to the rates calculated based on the 2018 ABS technology module: 1.4% and 5.8%, respectively (see Zolas et al., 2020b), even though there are some differences across the two modules in the way technologies are defined and the survey reference periods for measurement (2015-17 versus 2016-18).

[16] The 2019 ABS also queries firms on the intensity of use (low, moderate, or high use). The most heavily adopted technologies turn out to be also the most intensively used. Specialized software has the highest intensity of use with 44% of users reporting high and 35% reporting moderate use, followed by cloud and specialized equipment, each with 32%–33% high use and 36%–39% moderate use. AI and robotics have the lowest intensity of use with only 15%–18% of users reporting high use and 33%–35% moderate use.

Table 1: Technology adoption rates for processes and method and as part of goods and services, ABS data for 2016–2018.

	Technology users		Technology providers	
	Share of firms using technology (1)	Share of workers employed at firms using technology (2)	Share of firms selling technology (3)	Share of workers employed at firms selling technology (4)
Artificial Intelligence	3.2%	12.6%	0.5%	2.2%
Robotics	2.0%	15.7%	0.3%	1.8%
Dedicated Equipment	19.6%	36.4%	2.5%	4.8%
Specialized software	40.2%	64.4%	4.3%	7.8%
Cloud computing	34.0%	61.8%	3.5%	7.1%
Any technology	47.6%	69.9%	6.3%	11.1%

Notes: Data from the 2019 ABS technology module and authors' calculations. Technology use rates are based on firms' answer to questions E3 of the ABS: "During the three years 2016 to 2018, to what extent did this business use the following technologies in production processes for goods or services?" The table provides the share of firms that report using the technology in any capacity (low, moderate, or high use). Technology provision rates are based on firms' answer to questions E20 of the ABS: "During the three years 2016 to 2018, did this business sell the following technologies or goods or services that included the following technologies?" We exclude firms who either responded don't know or didn't respond to the question about technology use or provision in our calculations.

15.7% of US workers are employed at firms using robots. The same number is 12.6% for AI, 36.4% for dedicated equipment, 64.4% for specialized software, and 61.8% for cloud computing. These employment-weighted shares are the relevant object for gauging the importance of these technologies for labor and product markets.

A key advantage of the ABS module is that it distinguishes between *users* and *producers* of technology. Columns 1 and 2 reported the share of firms and workers at firms *using* these technologies as part of their production processes. But the ABS also asks firms whether they sold goods or provided services that embedded these technologies (e.g., whether the firm produces and sell robots or provides cloud-based solutions to customers). Column 3 shows that the supply side of these technologies is even more concentrated than their use, with only 0.3% of firms selling robots, 0.5% providing goods and services embedding AI algorithms, and 3.5% of firms selling cloud-based solutions. Column 4 shows that suppliers account for a small share of US employment too, so that much more workers are employed at firms using these technologies than at firms producing them. In the rest of this paper we will focus on users of advanced technologies.

Finally, the ABS also allow us to study whether these technologies are complementary, or whether they are being adopted by different firms. Table 2 shows that adopters tend to use multiple technologies at the same time, suggesting complementarities. For example, 86% (90%) of firms that use AI also use cloud (software), and 90% (88%) of firms that use robotics also use specialized software (equipment). Columns 4 and 5 of Table 2 show that this is particularly the case for cloud computing and specialized software, which are typically used to control and handle the computing needs of robotic systems and dedicated equipment.[17]

Table 2: Conditional adoption rates of multiple technologies, ABS data for 2016–2018.

	Share of firms using technology Y (column) conditional on using X (row)				
	Y =Artificial Intelligence (1)	Y =Robotics (2)	Y =Dedicated Equipment (3)	Y =Specialized Software (4)	Y =Cloud Computing (5)
X =Artificial Intelligence	100%	19.3%	54.4%	90.0%	85.8%
X =Robotics	30.2%	100%	87.7%	89.6%	73.1%
X =Dedicated Equipment	8.7%	9.0%	100%	82.9%	59.8%
X =Specialized Software	7.1%	4.5%	40.7%	100%	68.2%
X =Cloud Computing	7.9%	4.3%	34.4%	79.9%	100%
Unconditional rates	3.2%	2.0%	19.6%	40.2%	34.0%

Notes: Data from the 2019 ABS technology module and authors' calculations. The table reports the conditional probability of a firm using technology Y (reported across columns) given that it uses technology X (reported across rows). Technology use rates are based on firms' answer to questions E3 of the ABS: "During the three years 2016 to 2018, to what extent did this business use the following technologies in production processes for goods or services?" These conditional probabilities exclude firms who either responded don't know or didn't respond to the question about technology use.

4.2 Motivations for Technology Adoption

A common view is that advanced technologies facilitate the automation of tasks previously performed by labor. Industrial robots and dedicated equipment are being used in manufacturing to automate tasks such as welding, painting and assembly. AI is being used to create algorithms capable of achieving human proficiency at predictive tasks, such as controlling automated vehicles, trading, and medical diagnosing. And specialized software systems are capable of handling payrolls and sales. While these examples abound, we do not know the true extent to which firms are using these advanced technologies for automation. It could well be the case that firms use these technologies to control the quality of their processes, replace older vintages of machinery (instead of workers), or produce new tasks associated with an expanded variety of goods and services, alternative uses which do not involve the automation of tasks previously performed by labor. The distinction is consequential. As our theoretical framework demonstrates, the use of a technology for automation generates very different effects on firm and worker outcomes than other uses that make firms more productive without displacing workers from their tasks.

To understand the importance of automation and other uses of advanced technology, the survey asked adopters to identify their motivations for adoption from a list of six possibilities: i. to automate tasks performed by labor; ii. to upgrade outdated processes or methods; iii. to improve

[17] Likewise, some specialized software products such as Customer Relationship Management (CRM) or Enterprise Resource Planning (ERP) systems as well as some cloud services (e.g., AWS) have (or are planning to incorporate) some built-in AI capabilities. 59 percent of early AI adopters report using such tools to implement or test AI applications, according to a survey conducted by Deloitte (Loucks, Davenport and Schatsky, 2018).

the quality or reliability of processes; iv. to expand the range of goods and services provided; v. to adopt standards and accreditation; vi. some other reasons. Adopters were able to select all the motivations that applied, so responses are not exclusive.

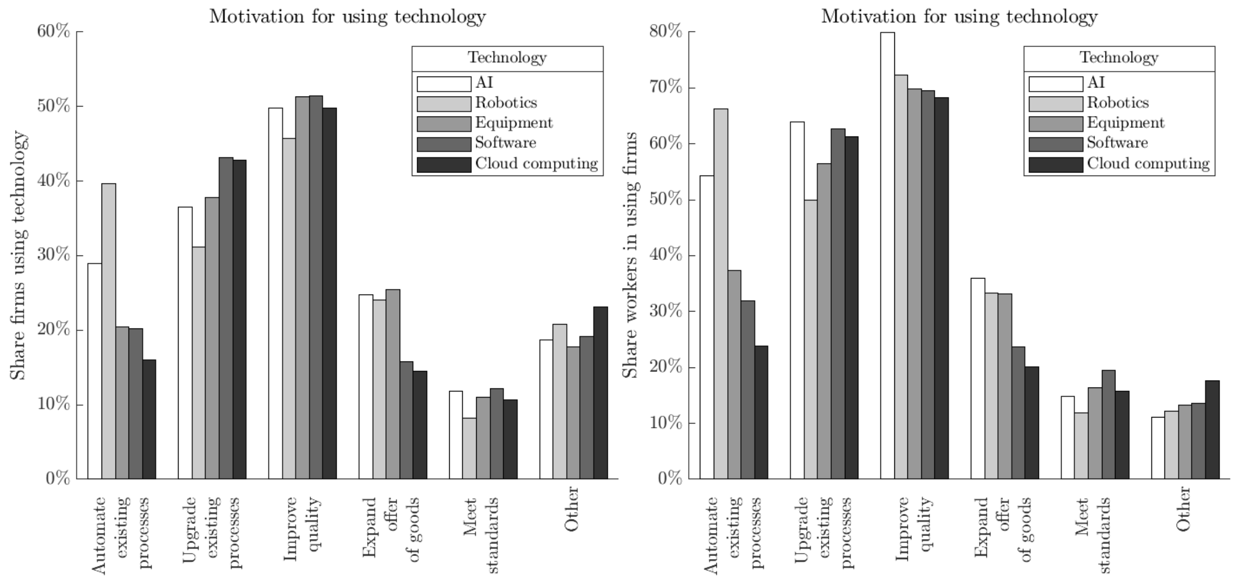

Panel A: Share of user firms

Panel B: Employment-weighted shares

Figure 1: Motivation for technology adoption, ABS data for 2016–2018. Tabulations based on questions E4 of the ABS, which asked firms "During the three years 2016 to 2018, why did this business adopt or use [each one of the five technologies surveyed in the ABS]? Select all that apply."

The share of adopting firms stating these motivations by technology is shown in Figure 1. Panel A reports the share of user firms reporting each motivation, while Panel B reports an employment-weighted share. In both panels, but especially in Panel B, we see that a significant share of adopters report using these technologies to automate tasks performed by labor, though this is not the main or only motivation. The estimates in Panel B show that AI and robotics are the two technologies with the greatest automation component, with 54% of firms using AI and 66% of the firms using robotics (in an employment-weighted sense) doing so to automate tasks. On the other hand, 37% of the firms using dedicated equipment and 32% of the firms using specialized software do so to automate tasks. Cloud computing is the technology that appears the least connected to automation, with 24% of firms using it to automate tasks.

The comparison between Panels A and B reveals that large firms are also more likely to use these technologies for automation conditional on their adoption. For example, the 40% of robotics-using firms adopting the technology for automation (Panel A) account for 66% of employment among these firms (Panel B).

Improving quality and the reliability of processes is the most common motivator. The employment weighted shares in Panel B show that 68%–80% of all users cite this as a motivation for using all technologies. About 50–64% of firms report using these technologies to upgrade outdated processes or methods, and 20–36% report using these technologies to expand the range of goods and services offered. Adopting standards and accreditation is the least common motivation for all technologies.[18]

In sum, the ABS points to automation being a distinct and important driver of the adoption of advanced technologies. Automation is as important as other uses leading to higher process quality and reliability, and more important than the use of these technologies associated with the introduction of new products. However, the ABS data also highlights that the extent to which advanced technologies are being used for automation varies with each technology, with AI and robotics being more closely linked to automation and software and cloud-computing systems having more diverse uses.

There are two interpretations of these findings. On the one hand, this could reflect some fundamental and permanent differences across technologies. For example, it might be the case that dedicated equipment and specialized software simply have more diverse applications, while most applications of robotics and AI entail the automation of some set of tasks. This would also explain the wider diffusion of software and dedicated equipment relative to AI and robotics. On the other hand, these differences might be temporary and simply reflect the maturity of these technologies. For example, it could be the case that technologies are initially deployed with an emphasis on automation, since these applications are more salient and easier to conceive. Over time, and as the technology matures and diffuses, new and more diverse applications of the technology emerge. While the cross-sectional ABS data does not allow us to tease apart these two explanations, we view this as an important question going forward.

4.3 The Exposure of US Workers to Automation

Using the ABS information on motivation, we can also compute the *exposure of US workers to automation*, defined as the share of workers employed at firms using advanced technologies to automate tasks. We see this measure as reflecting the importance of the use of advanced technologies for automation in US labor markets. Figure 2 reports our estimates of worker exposure to automation. 6.8% of US workers are employed at firms using AI for automation, and this number rises to 10.4% for robotics, 13.5% for dedicated equipment, 20.5% for specialized software, and 14.7% for cloud computing. Even though AI and robotics are more likely to be used to automate tasks, automation via dedicated equipment, specialized software, and cloud-based systems have been more important contributors on the aggregate due to their wider adoption and applicability.

The ABS data also allow us to gauge the importance of automation in manufacturing and outside of this sector. This is important because most studies have focused on automation via robotics in the manufacturing sector, and relatively less is known about the extent to which automation is also taking place in services and retail. Exposure to automation is higher in manufacturing, with 17.6%

[18] The general ranking of the motivations for technology adoption reported by firms matches the results from the 1991 Survey of Manufacturing Technology. The top benefit from technology use stated by plants in that survey was quality improvement, followed by labor cost reduction and flexibility increase (see Figure 3 in Dinlersoz and Wolf, 2018).

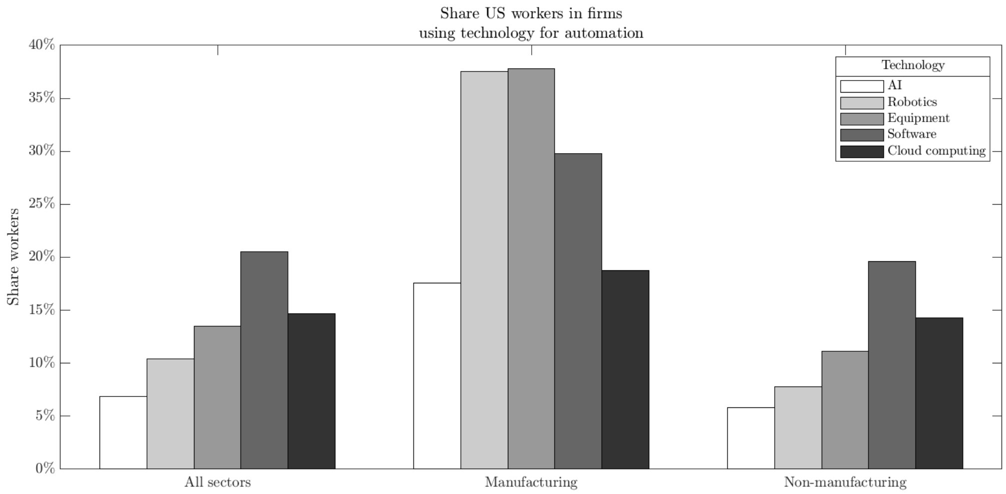

Figure 2: US workers' exposure to automation via advanced technologies, ABS data for 2016–2018. This exposure measure is computed as the share of the US workforce currently employed at firms using each technology for automation.

of US manufacturing workers employed at firms using AI for automation, and this number rising to 37.6% for robotics, 37.8% for dedicated equipment, and 29.8% for specialized software. However, automation is far from exclusive to manufacturing: 20% of US workers outside of manufacturing are employed at firms pursuing automation via specialized software systems.

While not all workers employed at firms using these technologies are (or were) subject to the effects of automation, these high exposure measures suggest that the use of advanced technologies for automation might be a relevant force affecting the US labor market, despite their limited diffusion across firms. This is because large firms are both more likely to adopt advanced technologies and to use these technologies for automation, and these firms account for a large share of labor demand.

4.4 Factors Adversely Affecting Technology Adoption

The ABS data reveal that a majority of firms, especially small firms, have not adopted advanced technologies. To understand the main factors limiting the adoption of advanced technologies, the 2019 ABS asked firms to identify all the factors limiting adoption from a list of 10 options (including costs, lack of data or skills, concerns about safety and regulations), or to indicate that no factors limited their adoption. Half of the firms that did not adopt technologies reported some factors that limited their adoption. As summarized in Panel A of Table 3, 45–50% of non-adopters (the vast majority of firms that reported some limiting factor) report that the advanced technologies in the ABS module are not applicable to their business. Conditional on the technology being applicable, the main adverse factor discouraging adoption is its high cost, with 7–9% of non users (about a fifth of the firms that reported some limiting factor) identifying high costs as the main bottleneck across all technology classes.

For firms that adopted the technology, the adverse factors listed may be interpreted as discour-

Table 3: Factors limiting the adoption of advanced technologies, ABS data for 2016–2018.

	Technology				
	Artificial Intelligence (1)	Robotics (2)	Dedicated Equipment (3)	Specialized Software (4)	Cloud Computing (5)
	Panel A: firms not using the technology				
No adverse factors	42%	41%	44%	44%	43%
Technology not applicable to this business	49%	50%	47%	46%	45%
Technology too expensive	7%	7%	7%	8%	7%
Technology not mature	2%	1%	0%	1%	1%
Lacked access to required data	1%	0%	0%	0%	1%
Required data not reliable	0%	0%	0%	0%	0%
Lacked access to human capital or talent	1%	1%	1%	1%	1%
Laws and regulations	1%	0%	0%	0%	0%
Concerns regarding safety and security	1%	0%	0%	1%	3%
Lacked access to capital	1%	1%	1%	1%	1%
	Panel B: firms using the technology				
No adverse factors	52%	64%	72%	77%	75%
Technology not applicable to this business	13%	8%	8%	7%	7%
Technology too expensive	17%	17%	12%	9%	6%
Technology not mature	10%	4%	1%	1%	2%
Lacked access to required data	4%	1%	1%	1%	1%
Required data not reliable	4%	2%	1%	1%	1%
Lacked access to human capital or talent	6%	4%	2%	2%	2%
Laws and regulations	4%	2%	2%	2%	2%
Concerns regarding safety and security	6%	3%	2%	3%	7%
Lacked access to capital	7%	7%	5%	3%	2%

Notes: Data from the 2019 ABS technology module based on authors' calculations. The table reports the share of non adopters (Panel a) and adopters (Panel b) that report each of the factors listed in the rows as adversely affecting their adoption of each technology, with separate technologies in different columns. The estimates reported above are based on responses to the following question in the 2019 ABS: "During the three years 2016 to 2018, indicate which factors adversely affected the adoption or utilization of the following technologies to produce goods or services. Select all that apply for each technology."

aging further adoption or the intensity of use. As summarized in Panel B, adopters faced fewer limitations as a whole, with 52–77% reporting no bottlenecks. On the other hand, some adopters identify lack of applicability (7–13% of firms) and high costs (6–17%) as the main factors limiting further adoption. The case of AI and robotics is particularly interesting, since these are the technologies with the lowest adoption rates and for which users reported the most bottlenecks (48% of AI users and 36% of robotics users reported some adverse factor that limited their adoption, compared to 25% of users for the remaining technologies). In contrast to the remaining technologies, users of AI and robotics see these technologies as lacking maturity, and identify the lack of human capital and financing as important bottlenecks for their adoption and more intense use.[19]

The findings from the ABS suggest that these advanced technologies have had limited applicability and require a high cost of adoption. This view aligns with our model, which sees advanced technologies as applicable to specific tasks—rather than general purpose technologies increasing the productivity of all firms at all industries irrespective of their task structure—and recognizes that there might be a high fixed cost of adoption. These two factors limit adoption but also imply that adoption concentrates in specific industries (those with the greatest applicability) and among large firms, a point that we discuss next.

5 Differences in Adoption Between and Within Industries

This section documents sizable differences in adoption rates between and within industries. Within an industry, adopters are larger, younger, have lower labor shares and higher labor productivity, and paid higher wages by 2015.

5.1 Differences in Adoption by Industry, Firm Size and Age

Figure 3 illustrates the differences in adoption rates both in terms of the share of firms using the technology (Panel A) and the share of workers at using firms (Panel B) across sectors. The adoption of advanced technologies is particularly high in manufacturing, though advanced technology use is far from exclusive to the manufacturing sector and is reasonably high in the information sector, professional services, healthcare, retail, and wholesale. The main exception to this pattern is robotics, which remains highly concentrated in manufacturing, with 8.7% of manufacturing firms using robotics and 45.1% of all manufacturing workers being exposed to this technology, while firms in other sectors exhibit much lower adoption rates.

The ABS data also reveals large variation across detailed industries in the same broad sectors. For example, in manufacturing, 87.4% of the workers in hardware manufacturing (NAICS code 3325), 81.8% of the workers in forging (NAICS 3321), 67.2% of the workers in motor vehicle manufacturing (NAICS 3363), and 75.4% of workers in dairy product manufacturing (NAICS 3115) are employed at firms using robotics, which are much higher than the manufacturing sector mean of 45.1%. One explanation for the importance of detailed industry in determining adoption patterns is that the nature of products and tasks varies across detailed industries, and this determines the applicability of advanced technologies to perform specific tasks in those industries.

5.2 Adoption Rates by Firm Size and Age

Even though industries play an important role in determining adoption, there is sizable variation in adoption across firms within detailed industries. Figure 4 explores the role of size and age, which

[19]In line with the ABS reports, a 2018 Deloitte survey of 1,100 US executives found that more than 30 percent of respondents reported challenges in implementation and integration, data issues (privacy as well as accessing data), cost, lack of skilled workers, and measuring/proving business value (see Loucks, Davenport and Schatsky, 2018).

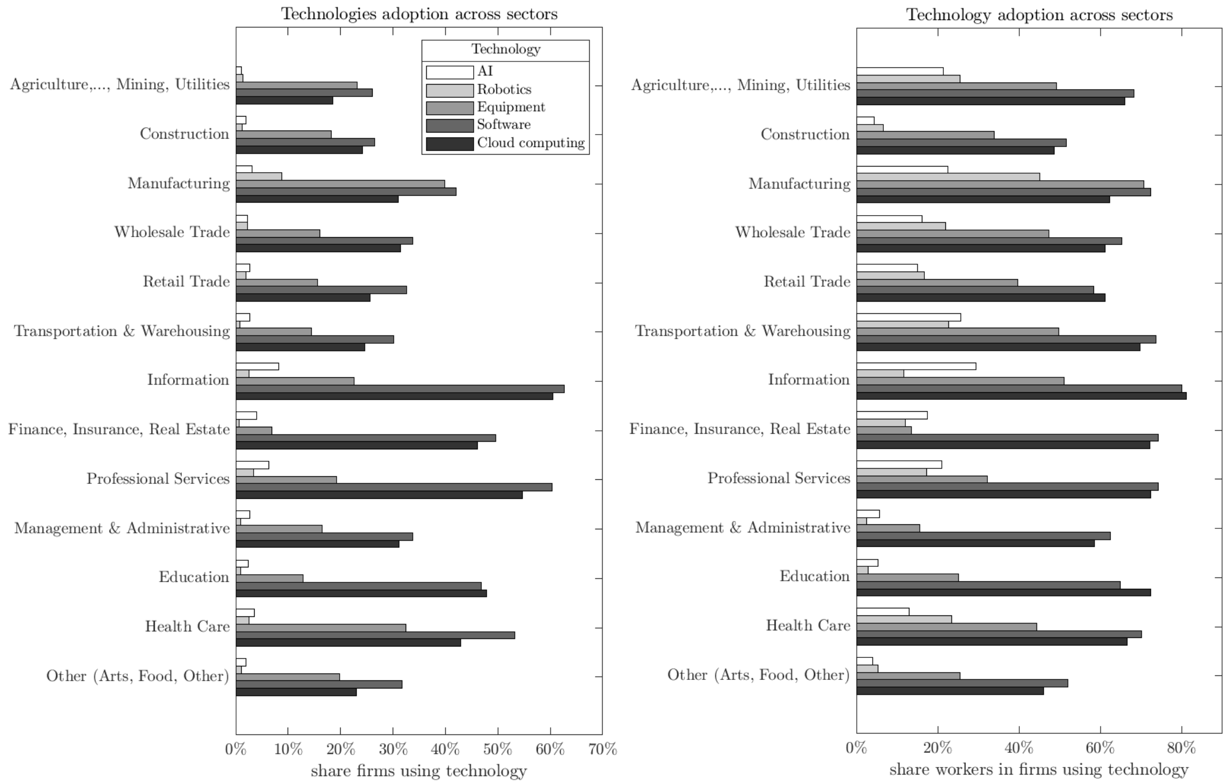

Panel A: Share of firms using technologies

Panel B: Share workers at firms using technologies

Figure 3: Technology use across US economic sectors, ABS data for 2016–2018. Technology use rates are based on firms' answer to questions E3 of the ABS, "During the three years 2016 to 2018, to what extent did this business use the following technologies in production processes for goods or services?" excluding firms who responded don't know or didn't respond.

our framework identifies as important dimensions mediating adoption decisions. We report for each technology the usage rates by 36 size–age categories, defined in terms of employment and age percentiles within each detailed 6-digit industry.[20] For example, Panel A provides the adoption rate of AI for firms belonging to the same size and age class in a detailed industry. The figure also reports the overall adoption rate for all firms of a size class.

Adoption rises significantly with size across all technologies. For example, 5.5% of firms in the top percentile of the employment distribution of their industries use AI, 5.1% use robots, 31.4% use dedicated equipment, 67.4% use specialized software, and 63.5% use cloud computing. In contrast, the adoption rate among firms between the 50th to 75th percentile of an industry employment distribution is much lower: 3.1% for AI, 1.7% for robots, 18.6% for dedicated equipment, 39.6% for specialized software, and 33.4% for cloud computing. These facts support the idea that automating tasks, or adopting advanced technologies more generally, involves large integration costs, which

[20] We assign firms to their predominant 6-digit NAICS industry in terms of payroll across all its establishments. In this and all subsequent exercises, the employment percentiles are defined based on the employment distribution from the LBD in each industry.

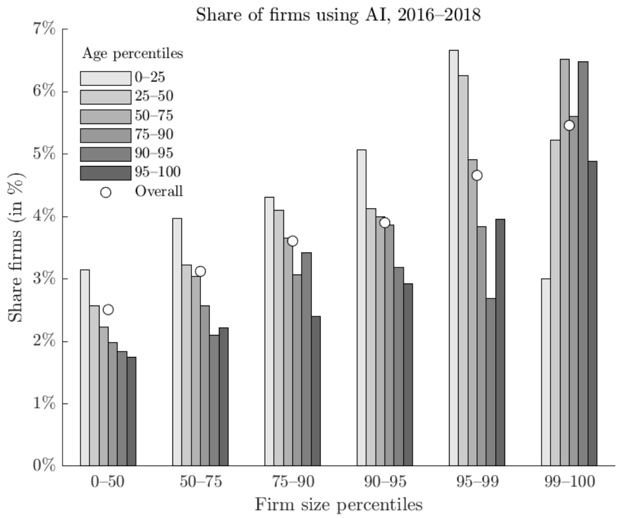

Panel A: Artificial intelligence

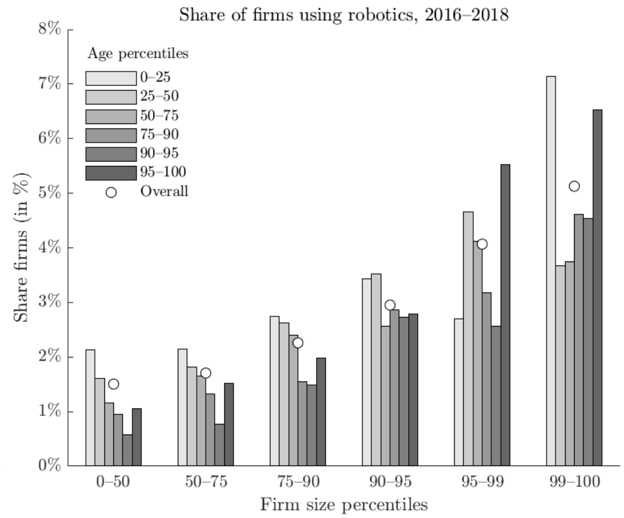

Panel B: Robotics

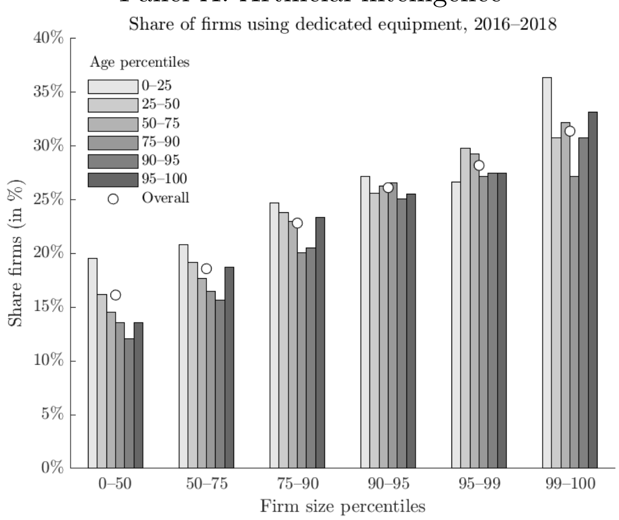

Panel C: Dedicated equipment

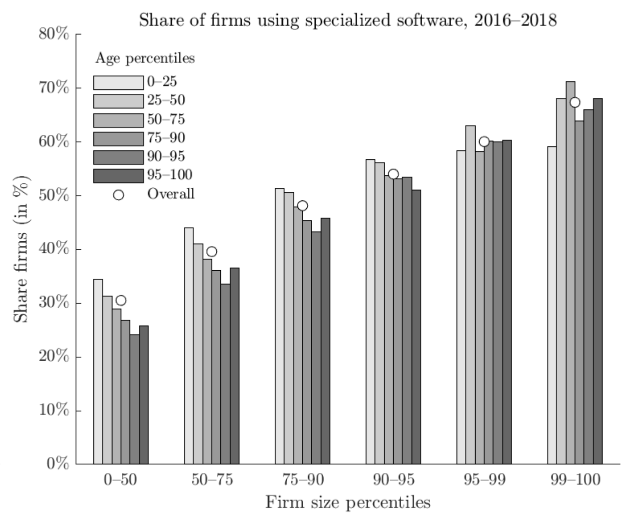

Panel D: Specialized software

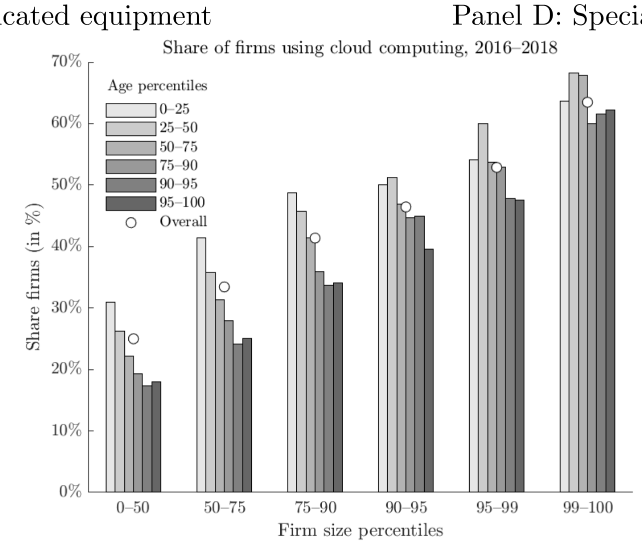

Panel E: Cloud computing

Figure 4: Adoption rates by firm size and age classes within detailed 6-digit industries, ABS data for 2016–2018.

imply that large and growing firms will select into adopting these technologies.[21]

Although not as strong as the pattern for size, we see that among firms of a given size, adoption tends to decrease with age. For most size classes, younger firms are more likely to adopt advanced technologies than older ones.[22] The declining adoption rates by age suggest that, as explained in the theory section, the fixed cost of adopting advanced automation technologies might be lower for younger firms, presumably because younger firms face fewer organizational barriers or do not have to pay the re-organizational costs of reallocating workers as they automate tasks. The patterns in Figure 4 also suggest that new entrants play an important role in the diffusion of advanced technologies, as is commonly assumed in models of technology diffusion (see, for example, Perla, Tonetti and Waugh, 2021; Hubmer and Restrepo, 2021). Likewise, these patterns point to the slowdown in entry as potentially contributing to the low adoption rates observed at the aggregate level, especially for smaller firms (see Decker et al., 2020, for evidence on the decline of entry and dynamism).

5.3 Correlates of Adoption

We now turn to a regression analysis to explore the correlates of technology adoption. We focus on a range of firm-level characteristics, in each case measured in 2015—a date that preceded the reference period of the ABS. Specifically, we estimate the following linear probability model for each technology:[23]

$$\text{Adopter}^{\text{tech}}_{f,i,s,a} = \alpha^{\text{tech}}_i + \gamma^{\text{tech}}_s + \lambda^{\text{tech}}_a + \beta^{\text{tech}} \cdot X_f + \varepsilon^{\text{tech}}_{f,i,s,a}. \tag{2}$$

The dependent variable is a dummy for whether firm f in industry i size class s and age a used each advanced technology during 2016–2018.[24] These models explain technology use as a function of 6-digit industry dummies α^{tech}_i, the size classes introduced before γ^{tech}_s and defined by employment percentiles in each detailed industry, and the age classes introduced before λ^{tech}_a and defined by age percentiles in each detailed industry. Finally, we also explore the role of other firm characteristics

[21] In line with this interpretation, Appendix C shows that adopters have had larger establishments than non-adopters from their same cohort at every age and that, for many cohorts, these differences preceded the arrival of advanced technologies. Moreover, these size differences between adopters and non-adopters have become smaller for more recent cohorts, presumably as these technologies become standardized and the costs of integration falls.

[22] One exception to this pattern is that technology adoption rates for firms in the 95+ age percentile group are higher than that of the 90–95 group. One difficulty in interpreting the results for the oldest firms is that the LBD only extends back to 1976, meaning that we cannot identify precise age values for firms born before 1977. Thus, the age distribution of firms in our ABS sample is truncated from above at 42 years, with all firms of age 42+ being assigned to the highest age percentile group.

[23] This timing choice does not imply any causality. First, adoption may have taken place prior to the reference period (the ABS only asked firms if they used the technology in 2016–2018, not if they adopted it during this period). Second, even firms that adopted during 2016–2018, might already be different in terms of unobservables by 2015. We interpret the estimates simply as identifying the main cross-sectional differences between users and other firms.

[24] We generate separate samples for each technology by dropping firms that answer either "don't know" or left blank questions regarding the intensity of their use of that technology (e.g. a firm that answers "don't know" to whether it uses robotics is not included in the sample for which we analyze robotics adoption). We also generate a separate sample for "any technology" by dropping firms that answer either "don't know" or left blank all technology use intensity questions.

from the LBD, including firms' labor productivity (revenue/employment), labor share of revenue (salaries and wages/revenue), and average wages paid in 2015 (salaries and wages/employment). In addition, all our specifications control for firms' employment shares across all US states, to account for their geographical location. We also include two measures that capture the industry diversification of firms and the importance of manufacturing as part of their activities. The first one is a dummy variable that equals 1 for firms that have some (but not all) of their establishments engaged in manufacturing activities. This dummy identifies firms with some manufacturing production. The second computes for these multi-sector firms the share of firm's employment classified in manufacturing based on its establishments' industry codes. All regressions use firm weights that are constructed to make the ABS sample representative of the universe of employer firms in the LBD.

Tables 4–5 present the estimates of equation (2), with a separate panel for each technology. Columns 1, 4, 7 in Table 4 (and columns 1 and 4 in table 5) show that, across all technologies there is a robust and sizable increase in the likelihood of technology usage with firm size. For example, an increase in within-industry size from the 0–50th employment percentile category to the 99+ percentile category is associated with an increase in adoption rate that ranges from 77% of the overall adoption rate (reported in the bottom rows) in the case of software to 220% in the case of robotics. Firm age, on the other hand, is negatively associated with technology usage, though these differences are not as large as those associated with size.

In line with the theoretical framework and the fact that a significant share of firms use advanced technologies for automation, these specifications also show that adopters have higher labor productivity. As shown in the model, labor productivity can be expressed as

$$\text{labor productivity} = \text{mean wage/firm labor share of revenue.}$$

Columns 2, 5, 8 in Table 4 (and columns 2 and 5 in table 5) exploit this decomposition and show that adopters higher labor productivity is driven both by higher average wages and a lower labor share of revenue. In particular, adopters pay higher wages and have lower labor shares than other firms of similar size and age in their same detailed industries.

Even though the correlations in Tables 4 do not provide guidance on the direction of causality, the fact that adopters paid higher wages is in line with our theory, which emphasizes the fact that higher wages generate incentives for automation.[25] Likewise, the fact that users of advanced technologies have lower labor shares and higher labor productivity is consistent with the use of these technologies for automation. Indeed our model clarifies that a low labor share and a high labor productivity are telltale signs of automation.

Columns 3, 6, 9 in Table 4 (and columns 3 and 6 in table 5) explore the role of industry di-

[25] A second interpretation of this finding which is also consistent with firms' assessments in Section 7 is that automation leads to a reallocation of labor from automated tasks to other roles, such as managerial, design, and engineering jobs that typically pay higher wages, which explains the higher mean wages. A third interpretation is that firms that adopt advanced technologies had a more skilled workforce to begin with, as in Doms, Dunne and Troske (1997).

Table 4: Regressions explaining the adoption of artificial intelligence, robotics, and dedicated equipment, ABS data for 2016–2018.

Dependent variable:	Artificial intelligence			Robotics			Dedicated equipment		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Employment percentile 50th-75th	0.009	0.008	0.009	0.008	0.008	0.008	0.060	0.059	0.060
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.004)	(0.004)	(0.004)
Employment percentile 75th-90th	0.013	0.013	0.013	0.016	0.016	0.015	0.096	0.094	0.095
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.005)	(0.005)	(0.005)
Employment percentile 90th-95th	0.017	0.017	0.017	0.026	0.026	0.025	0.114	0.111	0.113
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.008)	(0.008)	(0.008)
Employment percentile 95th-99th	0.017	0.017	0.017	0.029	0.029	0.027	0.128	0.126	0.126
	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)	(0.009)	(0.009)	(0.009)
Employment percentile 99th+	0.033	0.033	0.032	0.044	0.044	0.039	0.175	0.174	0.169
	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.021)	(0.021)	(0.022)
Age percentile 11th-50th	-0.005	-0.005	-0.005	-0.002	-0.002	-0.003	-0.014	-0.014	-0.014
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)	(0.005)
Age percentile 50th-75th	-0.009	-0.009	-0.009	-0.004	-0.004	-0.004	-0.034	-0.034	-0.034
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)	(0.005)
Age percentile 75th-90th	-0.012	-0.012	-0.012	-0.007	-0.007	-0.007	-0.039	-0.039	-0.039
	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)	(0.006)	(0.006)	(0.006)
Age percentile 90th-95th	-0.017	-0.017	-0.017	.	.	.	-0.049	-0.050	-0.049
	(0.003)	(0.003)	(0.003)				(0.009)	(0.009)	(0.009)
Age percentile 95th-99th	-0.021	-0.022	-0.021	-0.005	-0.005	-0.006	-0.068	-0.069	-0.068
	(0.005)	(0.005)	(0.005)	(0.002)	(0.002)	(0.002)	(0.012)	(0.012)	(0.012)
Age percentile 99th+	-0.010	-0.011	-0.011	.	.	.	-0.043	-0.044	-0.044
	(0.003)	(0.003)	(0.003)				(0.009)	(0.009)	(0.009)
Log of labor productivity	0.002		0.002	0.003		0.003	0.014		0.014
	(0.001)		(0.001)	(0.001)		(0.001)	(0.002)		(0.002)
Log of labor share		-0.001			-0.003			-0.011	
		(0.001)			(0.001)			(0.002)	
Log of average wage		0.003			0.003			0.017	
		(0.001)			(0.001)			(0.002)	
Multi-sector dummy			0.035			0.143			0.130
			(0.011)			(0.014)			(0.027)
Manufacturing employment share			-0.045			-0.020			0.153
			(0.029)			(0.047)			(0.103)
R-squared	1.6%	1.6%	1.6%	5.1%	5.1%	5.3%	11.8%	11.8%	11.8%
Mean adoption rate	0.032	0.032	0.032	0.02	0.02	0.02	0.196	0.196	0.196
Observations (rounded)	117,000	117,000	117,000	120,000	120,000	120,000	118,000	118,000	118,000

Notes: The table reports results from a regression of firm-level adoption on firm characteristics, 6-digit industry dummies, and employment shares by state. Columns 1–3 report results for the adoption of artificial intelligence. Columns 4–6 report results for the adoption of robotics. Columns 7–9 report results for the adoption of dedicated equipment. To protect confidentiality, in columns 4–6 the coefficients for firms in the age percentiles 75th to 90th and 90th to 95th are pooled together, as well as firms in the age percentiles 95th to 99th and above the 99th percentile. These coefficients are reported under the row corresponding to firms in the 75th to 90th and 95th to 99th age percentiles, respectively.

Table 5: Regressions explaining the adoption of specialized software and cloud computing, ABS data for 2016–2018.

Dependent variable:	Specialized software			Cloud computing		
	(1)	(2)	(3)	(4)	(5)	(6)
Employment percentile 50th-75th	0.121	0.116	0.121	0.110	0.104	0.110
	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)
Employment percentile 75th-90th	0.186	0.178	0.185	0.171	0.163	0.171
	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)
Employment percentile 90th-95th	0.240	0.231	0.240	0.228	0.218	0.228
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Employment percentile 95th-99th	0.272	0.262	0.272	0.278	0.267	0.277
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Employment percentile 99th+	0.320	0.316	0.319	0.338	0.334	0.337
	(0.022)	(0.022)	(0.022)	(0.021)	(0.022)	(0.022)
Age percentile 11th-50th	-0.024	-0.025	-0.024	-0.044	-0.045	-0.044
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Age percentile 50th-75th	-0.051	-0.052	-0.051	-0.087	-0.089	-0.087
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Age percentile 75th-90th	-0.056	-0.059	-0.056	-0.112	-0.115	-0.112
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Age percentile 90th-95th	-0.076	-0.080	-0.076	-0.129	-0.133	-0.129
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Age percentile 95th-99th	-0.116	-0.120	-0.116	-0.178	-0.183	-0.178
	(0.015)	(0.015)	(0.015)	(0.014)	(0.014)	(0.014)
Age percentile 99th+	-0.070	-0.074	-0.071	-0.119	-0.123	-0.119
	(0.010)	(0.010)	(0.010)	(0.009)	(0.009)	(0.009)
Log of labor productivity	0.040		0.040	0.043		0.043
	(0.002)		(0.002)	(0.002)		(0.002)
Log of labor share		-0.028			-0.030	
		(0.002)			(0.002)	
Log of average wage		0.056			0.060	
		(0.003)			(0.002)	
Multi-sector dummy			0.041			0.048
			(0.030)			(0.025)
Manufacturing employment share			-0.204			-0.337
			(0.131)			(0.110)
R-squared	14%	14.1%	14%	14.2%	14.4%	14.2%
Mean adoption rate	0.402	0.402	0.402	0.34	0.34	0.34
Observations (rounded)	117,000	117,000	117,000	118,000	118,000	118,000

Notes: The table reports results from a regression of firm-level adoption on firm characteristics, 6-digit industry dummies, and employment shares by state. Columns 1–3 report results for the adoption of specialized software. Columns 4–6 report results for the adoption of cloud computing.

versification and the importance of manufacturing in firms' activities. While adopters tend to be diversified and have some manufacturing activities, their share of manufacturing is lower than other multi-industry firms (with the exception of dedicated equipment for which the association is positive), though this relationship is not precisely estimated. Through the lens of the framework, this is consistent with adopters automating production roles for labor and reallocating their workforce towards design and non-manufacturing roles.[26]

Although not reported, all regressions control for detailed industry dummies and state employ-

[26]Recent research has found that continuing manufacturing firms have greatly increased their employment of non-manufacturing workers, contributing to the process of structural transformation in the economy (see Fort, Pierce and Schott, 2018). The findings here suggest that the adoption of advanced technologies might facilitate this process taking place within firms.

ment shares. Despite the fact that we have this rich set of covariates, the specifications generate low R^2 values, ranging from 1.6% in the case of AI to 14% in the case of cloud and specialized software, pointing to a large role for unobserved idiosyncratic factors at the firm level.[27] In terms of the explanatory power of the covariates, we find that 6-digit NAICS industry dummies account for 66%–88% of the explained variation in adoption; while firm size classes explain 6.3%–23%. On the other hand, the geographic location of a firm plays a small role, which suggest that firms in the same detailed industry and of similar size tend to have similar adoption rates across all US locations.

The appendix provides a series of related exercises. Tables A-1–A-2 repeat the regression analysis for manufacturing and non-manufacturing industries separately, uncovering similar patterns across sectors.[28] Finally, Table A-3 repeats our analysis but with the dependent variable now being an indicator of whether the firm uses the technology for automation (as opposed to not using it at all). Here too, we find very similar results.

6 Advanced Technologies and Labor Productivity

We now explore the link between technology adoption and labor productivity (measured as revenue per worker) in more detail. As noted in Section 3, adopters should achieve higher labor productivity for two reasons. First, when used for automation, advanced technologies allow firms to produce in a more capital-intensive way, by relying more on specialized equipment and software and less on labor. Second, these technologies may lead to lower employment of less-skilled workers and increase the hiring of more-skilled workers, and this effect on skill composition can also increase labor productivity. This is very different from a factor-neutral increase in productivity, which increases TFP but does not alter labor productivity, since firms expand their sales and employment by the same amount.

Despite this clear theoretical prediction, the adoption of the advanced technologies surveyed in the ABS took place against the backdrop of slowing aggregate productivity growth in the US, raising questions about whether new technologies are indeed increasing labor productivity. We document that despite the aggregate trends, the adoption of advanced technologies has been associated with rising labor productivity for adopting firms, both inside and outside of manufacturing.

To explore this in detail, we estimate regressions of the form

$$\text{Log labor productivity}_{f,i,s,a} = \alpha_i + \gamma_s + \lambda_a + \beta \cdot \text{Technology Adoption}_f + e_{f,i,s,a}. \tag{3}$$

This regression explains labor productivity, measured as revenue per worker in 2019, as a function of detailed industry fixed effects α_i, size-class dummies defined in terms of employment percentiles γ_s, age classes λ_a and different measures of technology adoption from the 2019 ABS. As before

[27] This relatively low explanatory power of observable firm characteristics also emerges in the analysis of AI adoption based on a set of AI-related technologies in the 2018 ABS technology module (see McElheran et al., 2022).

[28] To protect confidentiality, the oldest age group now contains firms in the 96+ percentiles of the age distribution within an industry, rather than the 99+ category used in Tables 4–5 for the general samples that include all sectors.

we interpret this regression as descriptive and emphasize that this approach does not necessarily recovers the causal effect of adoption on labor productivity.

Table 6 presents our estimates of equation (3). Columns 1–4 provide estimates for all sectors; while columns 5–7 and 8–10 provide separate estimates for manufacturing and non-manufacturing. The specification in column 1 contains only the size and age percentile categories. This specification shows that, starting at the median firm, there is an increasing relationship between firm size and labor productivity, with the largest firms in each 6-digit industry achieving 26.3% higher labor productivity than the middle firms in their industry (those in the 50th–75th employment percentiles).[29] This is the *superstar firm* phenomenon: large firms that manage to have high sales without necessarily hiring more workers than their competitors, documented before in work by Autor et al. (2020) and Kehrig and Vincent (2021). Columns 5 and 8 show that a similar phenomenon emerges in manufacturing and non-manufacturing. In addition, labor productivity declines with age (with the exception of the oldest firm group). Although not reported, these specifications also include detailed 6-digit industry dummies, which shows that the superstar phenomenon is visible even within detailed industries.[30]

Columns 2, 6, 9 explore the contribution of advanced technologies to the observed differences in labor productivity by including a dummy for firms that adopted any of the advanced technologies in the ABS. Column 3, 7, 10 go one step further and include separate dummies for each of the five technologies in the ABS. Across these specifications, we see that technology adoption is associated with a 14.2% increase in labor productivity (11.3% in manufacturing and 14.3% outside of manufacturing). Cloud, robotics, and specialized software have positive and significant association with labor productivity, whereas AI and specialized equipment are not highly correlated with productivity, though the individual coefficients are hard to interpret given that many of these technologies are adopted jointly.[31] Finally, we find in column 4 that the number of technologies that a firm uses is associated positively with its productivity. Firms with a single technology have a 9.1% higher labor productivity than firms with none, and at the extreme, firms with all five technologies have a 25.8% higher labor productivity.[32]

[29] One concern with the regression specifications for labor productivity is division bias, as firm size appears in the denominator of the dependent variable in addition to being an explanatory variable (as firm size percentile categories). To alleviate this concern, Table A-4 in Appendix B repeats the analysis using one-year lagged firm size to form the size percentile categories. The general pattern for the estimated firm size and age coefficients is robust to this alternative measure of size. The coefficients for technology indicators attenuate slightly, but still indicate a positive association with labor productivity for cloud, robots, and software, and for the presence of any technology in general.

[30] Industry dummies are the most important factor accounting for differences in labor productivity in our analysis, accounting for more than 80% of the explained variation in all cases. This is not surprising, given the large differences in intermediate input use across industries. Foster, Haltiwanger and Krizan (2001) present evidence that within detailed industries there is a strong positive correlation between measured TFP, measured value added per worker and gross output per worker across businesses.

[31] We also find that the adoption of any technology motivated by automation is associated with a 15.7% higher labor productivity; whereas firms that adopt the technologies but do not report doing this for the purposes of automating tasks achieve a 13.8% higher labor productivity. This suggest that firms using these advanced technologies explicitly for automation get a larger increase in their labor productivity, though this difference is only marginally significant from an statistical viewpoint.

[32] These findings are consistent with earlier studies that also indicate a positive connection between labor productivity and technology use using the SMT (see Dinlersoz and Wolf, 2018).

Table 6: Regressions explaining firm labor productivity in 2019 as a function of technology adoption, ABS data for 2016–2018.

	Dependent Variable: Log of Labor Productivity in 2019									
Sector:	All sectors combined				Manufacturing			Non-manufacturing		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Employment percentile 0th–50th	0.178	0.194	0.198	0.197	0.008	0.025	0.032	0.186	0.202	0.207
	(0.008)	(0.008)	(0.008)	(0.008)	(0.014)	(0.014)	(0.014)	(0.008)	(0.008)	(0.008)
Employment percentile 75th–90th	0.025	0.016	0.013	0.014	0.142	0.129	0.121	0.019	0.011	0.008
	(0.010)	(0.009)	(0.009)	(0.009)	(0.020)	(0.020)	(0.021)	(0.010)	(0.010)	(0.010)
Employment percentile 90th–95th	0.065	0.049	0.041	0.045	0.309	0.291	0.277	0.055	0.039	0.032
	(0.014)	(0.014)	(0.014)	(0.014)	(0.022)	(0.022)	(0.022)	(0.015)	(0.015)	(0.015)
Employment percentile 95th–99th	0.116	0.094	0.083	0.088	0.437	0.412	0.384	0.102	0.080	0.070
	(0.017)	(0.017)	(0.017)	(0.017)	(0.030)	(0.030)	(0.030)	(0.018)	(0.018)	(0.018)
Employment percentile 99th+	0.263	0.235	0.219	0.227	0.584	0.554	0.504	0.253	0.225	0.210
	(0.040)	(0.039)	(0.039)	(0.039)	(0.040)	(0.041)	(0.041)	(0.041)	(0.041)	(0.041)
Age percentile 10th–50th	0.007	0.011	0.012	0.012	-0.040	-0.033	-0.031	0.009	0.013	0.014
	(0.016)	(0.015)	(0.015)	(0.015)	(0.033)	(0.033)	(0.033)	(0.016)	(0.016)	(0.016)
Age percentile 50th–75th	-0.022	-0.013	-0.010	-0.011	-0.099	-0.089	-0.085	-0.020	-0.011	-0.008
	(0.016)	(0.016)	(0.016)	(0.016)	(0.033)	(0.033)	(0.033)	(0.017)	(0.017)	(0.017)
Age percentile 75th–90th	-0.040	-0.028	-0.024	-0.026	-0.117	-0.102	-0.098	-0.037	-0.026	-0.022
	(0.017)	(0.017)	(0.017)	(0.017)	(0.036)	(0.036)	(0.036)	(0.017)	(0.017)	(0.018)
Age percentile 90th–95th	-0.072	-0.059	-0.053	-0.056	0.045	0.070	0.084	-0.071	-0.058	-0.052
	(0.022)	(0.022)	(0.022)	(0.022)	(0.066)	(0.067)	(0.067)	(0.023)	(0.023)	(0.023)
Age percentile 95th–99th	-0.123	-0.104	-0.094	-0.098	0.161	0.183	0.186	-0.120	-0.100	-0.091
	(0.032)	(0.032)	(0.032)	(0.032)	(0.171)	(0.171)	(0.168)	(0.033)	(0.032)	(0.033)
Age percentile 99th+	-0.025	-0.012	-0.008	-0.010	-0.128	-0.116	-0.114	-0.025	-0.012	-0.007
	(0.021)	(0.021)	(0.021)	(0.021)	(0.037)	(0.037)	(0.037)	(0.022)	(0.022)	(0.022)
Technology user		0.142				0.113			0.143	
		(0.007)				(0.014)			(0.008)	
Artificial intelligence			-0.003				0.054			-0.003
			(0.019)				(0.041)			(0.019)
Cloud computing			0.111				0.051			0.115
			(0.009)				(0.017)			(0.009)
Robotics			0.098				0.111			0.066
			(0.021)				(0.022)			(0.026)
Specialized software			0.084				0.085			0.083
			(0.009)				(0.021)			(0.010)
Dedicated equipment			-0.018				-0.002			-0.021
			(0.010)				(0.019)			(0.010)
One technology				0.091						
				(0.011)						
Two technologies				0.157						
				(0.009)						
Three technologies				0.174						
				(0.011)						
Four technologies				0.216						
				(0.021)						
Five technologies				0.258						
				(0.043)						
R-squared	30.8%	31.2%	31.2%	31.3%	17.9%	18.3%	18.6%	31.5%	31.9%	32.0%
Observations (rounded)	101,000	101,000	101,000	101,000	20,000	20,000	20,000	81,500	81,500	81,500

Notes: The table reports results from a regression of firm labor productivity on size and age groups, 6-digit industry dummies, employment shares by state, and measures of technology adoption from the 2019 ABS. Columns 1–4 report results for all sectors. Columns 5–7 report results for manufacturing firms. Columns 8–10 report results for firms outside of manufacturing. To protect confidentiality, the number of observations has been rounded.

The estimates from column 3 imply that, from a descriptive viewpoint, the uneven adoption of advanced technologies by large firms explains close to a third of the superstar phenomenon. That is, conditional on the technology measures from the ABS, we see that the largest firms in each industry have a 21.9% higher labor productivity than mid-sized firms (as opposed to 26.3% when technology is not accounted for in Column 1). Thus, technology explains 20% of the productivity gap between firms above the 99th employment percentile relative to mid-size firms. Likewise, technology explains 33% of the labor productivity gap between firms in the 96th to 99th employment percentile relative to mid-size firms, and 46% of the labor productivity gap between firms in the 91th to 95th employment percentile relative to mid-size firms. This is also the case inside and outside of manufacturing, where differences in technology adoption account for 20% to 30% of the superstar phenomenon in a descriptive sense.

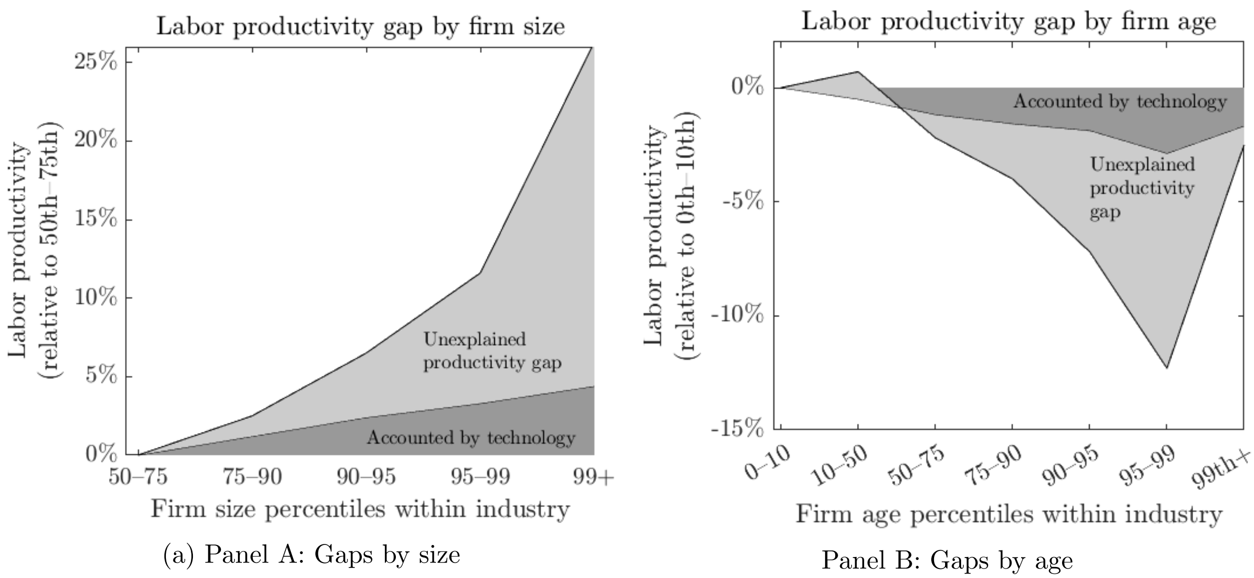

Figure 5: Estimates of differences in labor productivity gaps by size and age and the share of these differences explained by the adoption of advanced technologies, as measured in the 2019 ABS.

Likewise, we find that differences in technology adoption between young and older firms account for 26% of the labor productivity gap between firms in the 96th to 99th age percentiles and the youngest firms in the 0th to 10th age percentiles of each detailed industry. Figure 5 illustrates these finding graphically by plotting the labor productivity gaps as a function of firm-size class and age, indicating the share of these gaps explained by the measures of technology adoption from the ABS and the part left unexplained.

7 Implications for Demand for Workers and Skills

This section summarizes firms' assessments on the impact of advanced technologies on their workforce. The ABS asked adopting firms to assess the impact that advanced technologies have had on the employment level. Respondents were given 3 options: "Increase", "Decrease", and "Did Not Change". Figure 6 plots the self-reported direction of employment changes separately for each technology. The figure provides the share of firms (both the simple share and an employment-weighted version) reporting an increase or a decrease in employment (with the share of firms reporting no change given by the complement of these two). For simplicity, we focus on the employment weighted shares in our discussion. Across all technologies, most firms claim that the use of advanced technologies did not change their employment levels in recent years, with 67–78% of firms selecting this response.

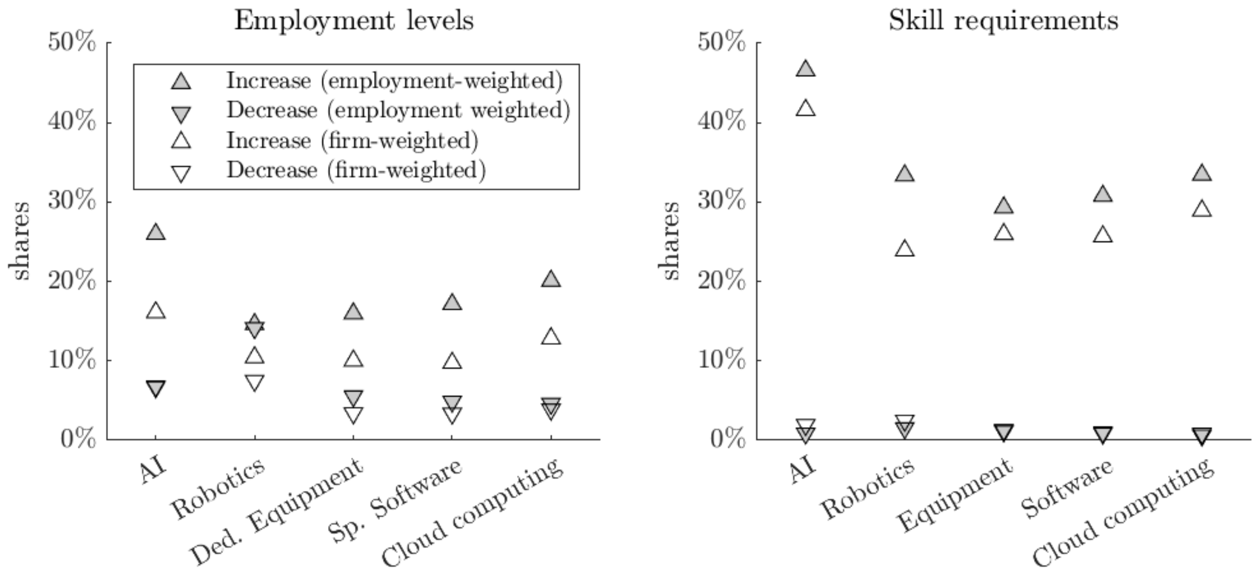

Figure 6: Reported changes in employment levels and skill demand by firms adopting advanced technologies, employment-weighted shares from 2019 ABS.

A small share of firms report positive or negative employment changes caused by the adoption of the technologies in the ABS. The share of employment-weighted firms reporting an increase in employment is of 26% for users of AI; while robotics is the technology most closely associated with employment decreases, with the share of employment-weighted firms reporting a decrease in employment due to the use of robots rising to 14% of firms (roughly the same share reporting an increase in employment attributed to this technology).[33] One caveat is that these responses reflect self-assessments by firms, and some firms may be particularly reluctant to divulge information on workforce reductions through technology adoption; while other firms might default to reporting no employment changes.

We interpret firms' assessments as pointing to limited and ambiguous employment effects of advanced technologies at the firm level. This finding aligns with the theoretical framework, which

[33]Figures A-1 and A-2 break down these answers by firm size and age. The 2019 ABS technology module also asked firms to assess the effect of technology on the number of different types of workers, including: production workers, non-production workers, supervisory workers, and non-supervisory workers. Here too, most firms report no changes in employment levels for these workers. One notable case is that of robotics, where the employment-weighted share of firms reporting a decrease in the employment of production workers exceeds the share reporting an increase.

highlights the fact that automation will have an ambiguous effect on firm employment. The framework also clarifies that it is reasonable to expect automation to increase employment in some firms while at the same time it reduces employment in others, as we find in the data. The possibility that advanced technologies have a limited effect on employment also underscores the importance of the displacement effects from automation. Consider, for example, a technological improvement that increases productivity in a factor neutral way, such as improvements in product quality, but do not involve the automation of tasks performed by labor. Our model shows that these factor neutral technological developments should always increase sales and employment proportionally.[34]

The 2019 ABS technology module also provides information on changes in the demand for skills linked to the adoption of advanced technologies. Firms were asked whether the skill level of their workers changed as a result of technology use, with the response options of "Increase", "Decrease", and "No Change". Panel B in Figure 6 plots the share of reported changes in skill attributed to technology adoption (again, weighted by employment). We now see a more sizable share of firms reporting an increase in their demand for skills, ranging from 30% to 50% of users in an employment-weighted sense. In addition, we see almost no firm reporting a reduction in their demand for skills.

The results from firms' assessments indicate that advanced technologies and their use for automation has resulted not so much in changes in firm-level employment but in changes in their employment composition, with firms increasing their demand for skills. This finding is in line with the theory, which suggests that the use of advanced technologies involves a reassignment of labor from automated tasks to other complementary roles, including the maintenance, programming, and operation of specialized machinery. Firms' responses also align with recent work highlighting the fact that the adoption of advanced technologies and robots is associated with significant changes in the workforce composition of firms, measured in terms of their occupational structure or the skill level of their workers (Dinlersoz and Wolf, 2018; Humlum, 2020; Bonfiglioli et al., 2020; Rodrigo, 2021, see, for example,).[35] The relatively high incidence of skill upgrading reported by firms suggests that the use of advanced technologies might be an important force contributing to the observed changes in the occupational and wage structure of the US economy over the last 40 years, though quantifying the contribution of these technologies to these shifts is beyond the scope of our paper (see Acemoglu and Restrepo, 2021b, for more on this quetsion).

[34]In the simple model of Section 3, this is always the case since firms face a constant elasticity of demand and have a constant markup, which makes their passthrough of marginal cost to price equal to 1. In practice, firms might expand their sales more than employment if their passthrough of marginal cost to price lies below 1. But for most reasonable values of the demand elasticity and passthroughs, we should see an expansion of both employment and sales in response to higher firm TFP.

[35]In contrast, previous work argued that the presence of skilled workers allowed firms to adopt advanced technologies, but that the extent of skill upgrading in response to the adoption of these technologies was limited. For example, Doms, Dunne and Troske (1997) conduct a longitudinal analysis of plants in the US based on the Survey of Manufacturing Technology, which indicates that plants with higher worker skills have higher technology adoption, but there is little skill upgrading in response to technology use.

8 Conclusion

A lack of comprehensive data at the firm level has precluded a detailed assessment of the current state of advanced technology use by US firms and these technologies' impact on productivity and the workforce. Recent surveys conducted by the Census Bureau as part of the 2018 and 2019 ABS fill this gap and offer new insights. Using the data collected by the technology module included in the 2019 ABS, we have provided new measures of the diffusion of five key technologies—AI, robotics, dedicated equipment, specialized software, and cloud—and documented the relationship between their adoption and firm characteristics and workforce outcomes. While these technologies (especially AI and robotics) have low diffusion rates among firms, a significant fraction of the US workforce are employed in firms using these technologies, because larger firms are much more likely to adopt them.

We documented a number of descriptive facts, which are mostly novel and complement previous work by Zolas et al. (2020a). Most importantly:

1. While the adoption of advanced technologies remains low for the period 2016–2018, with only 2% of firms currently using robotics as part of their processes and methods and 3.2% using AI, a significant fraction of workers are exposed to these technologies: 12.6% of the US workforce is employed in firms using AI technologies between 2016-2018, and the the analogous number is 15.7% for robotics, 64.4% for specialized software, 36.4% for dedicated equipment, and 61.8% for cloud computing. In manufacturing, worker exposure to advanced technologies is even higher: 23% for AI, 45% for robotics, 71% for dedicated equipment, and 72% for specialized software.

2. Automation is not the only objective of the adoption of these advanced technologies, but it is an important motivation for the adoption of AI and robotics, and to a lesser extent for dedicated equipment and specialized software.

3. In addition to larger firms, younger firms are much more likely to adopt advanced technologies than older firms. More productive, lower labor share, and higher-wage firms are also more likely to adopt these technologies, which is in line with our conceptual framework. There is also considerable variation between industries, even within the manufacturing sector.

4. We show that adopting firms have significantly higher labor productivity, thus providing the first comprehensive evidence on the relationship between various types of advanced technologies and productivity at the firm level for the US.

5. Firms' self-assessments generally point to an increase in the relative demand for skill but limited or ambiguous effects on their employment level. Combined with the theoretical expectation that firms adopting advanced technologies reduce their production costs and expand at the expense of their rivals in Koch, Manuylov and Smolka (2021) and Acemoglu, Lelarge

and Restrepo (2020), this evidence weighs against the view that advanced technology adoption will lead to higher employment, especially for low-skill workers.

This paper has focused on presenting a new set of descriptive facts. Moving forward, there are several interesting directions for future research. Many of these directions will also benefit from future planned ABS modules, which will add a longitudinal dimension to the data set. Here we list some of these directions:

- Future work can explore both whether the correlation between advanced technologies and labor productivity is causal at the firm level and how it aggregates to the industry and the economy. Composition effects and impacts of new technologies on markups will be particularly important in understanding these implications.

- Industry-level and aggregate employment implications of new technologies needs further study as well. To do this, one can estimate the impact of advanced technologies not just on adopting firms' employment and skill demand, but on their rivals. If effects on rivals are negative and large, advanced technologies can have negative consequences, and whether they do or not and how this varies across different classes of technologies are central questions for future research.

- It would also be interesting to explore how advanced technologies impact the economy by expanding the range of goods and services and enabling quality upgrades. The ABS points to automation being an important driver of the adoption of advanced technologies, with automation being as important as expanding the range of goods and services offered by firms in driving adoption. However, the ABS data also highlights that the extent to which advanced technologies are being used for automation varies with each technology and across firms, with a sizable share of firms reporting not using these technologies for automation. Understanding the determinants of these different motivations and uses is a fruitful area of research.

- Another important area is the study of whether labor shortages, for example such as those caused by the COVID-19 pandemic, trigger further automation and how permanent such shortage-induced adoption decisions will be.[36]

References

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2020): "AI and Jobs: Evidence from Online Vacancies," Working Paper 28257, National Bureau of Economic Research.

[36]Examples of research predicting an increase in investment in automation technologies in response to the Covid pandemic shock and forecasting related impacts include Autor and Reynolds (2020) and Chernoff and Warman (2020). Data enabling direct measurement of the association between the pandemic and technology adoption is still rare. Exceptions include Comin et al. (2022) (conducting a survey of firms in three developing countries and finding that pre-COVID digital technology adopters were better able to weather the pandemic and more likely to increase use of digital technologies than less advanced firms) and Alekseev et al. (2020) (conducting a survey of small businesses and finding that about one-half dealt with the pandemic by providing online services and about one-third expanded digital payments). The US Census Bureau added a question to Phase 7 (November 2021-January 2022) of its Small Business Pulse Survey on whether firms changed any of eight business practices since the onset of the pandemic, including adoption and expansion of digital technologies Bureau (2021).

Acemoglu, D., C. Lelarge, and P. Restrepo (2020): "Competing with Robots: Firm-Level Evidence from France," in *AEA Papers and Proceedings*, vol. 110, pp. 383–88.

Acemoglu, D., and P. Restrepo (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128(6), 2188–2244.

——— (2021a): "Demographics and Automation," *The Review of Economic Studies*, 89(1), 1–44.

——— (2021b): "Tasks, Automation, and the Rise in US Wage Inequality," Discussion paper, National Bureau of Economic Research. Paper No 28920.

Aghion, P., C. Antonin, S. Bunel, and X. Jaravel (2020): "What are the labor and product market effects of automation? new evidence from france," Discussion paper, CEPR Discussion Paper No. DP14443.

Alekseev, G., S. Amer, M. Gopal, T. Kuchler, J. Schneider, J. Stroebel, and N. C. Wernerfelt (2020): "The Effects of COVID-19 on U.S. Small Businesses: Evidence from Owners, Managers, and Employees," Working Paper 27833, National Bureau of Economic Research.

Alekseeva, L., J. Azar, M. Giné, S. Samila, and B. Taska (2021): "The demand for AI skills in the labor market," *Labour Economics*, 71, 102002.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020): "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 135(2), 645–709.

Autor, D., and E. Reynolds (2020): "The Nature of Work after the COVID Crisis: Too Few Low-Wage Jobs," Discussion paper, The Hamilton Project.

Babina, T., A. Fedyk, A. He, and J. Hodson (2021): "Artificial Intelligence, Firm Growth, and Industry Concentration," Discussion paper, Mimeo, Columbia GSB.

Basker, E., L. Foster, and S. Klimek (2017): "Customer-employee substitution: Evidence from gasoline stations," *Journal of Economics & Management Strategy*, 26(4), 876–896.

Bessen, J., M. Goos, A. Salomons, and W. Van den Berge (2019): "Automatic reaction-what happens to workers at firms that automate?," Discussion paper, Mimeo, Boston University School of Law.

Bloom, N., T. A. Hassan, A. Kalyani, J. Lerner, and A. Tahoun (2021): "The diffusion of disruptive technologies," Discussion paper, National Bureau of Economic Research.

Bonfiglioli, A., R. Crinò, H. Fadinger, and G. Gancia (2020): "Robot Imports and Firm-Level Outcomes," Discussion paper, CEPR Discussion Papers No. 14593.

Brynjolfsson, E., W. Jin, and K. McElheran (2021): "The power of prediction: predictive analytics, workplace complements, and business performance," *Business Economics*, 56(4), 217–239.

Brynjolfsson, E., and A. McAfee (2014): *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. Norton & Company.

Bureau, U. C. (2021): "Questionnaire for Small Business Pulse Survey, Phase 7," .

Calvino, F., S. DeSantis, I. Desnoyers-James, S. Formai, I. Goretti, S. Lombardi, F. Manaresi, and G. Perani (2022): "Closing the Italian digital gap: The role of skills, intangibles and policies," OECD Science, Technology and Industry Policy Papers 126, OECD Publishing.

Cheng, H., L. A. Drozd, R. Giri, M. Taschereau-Dumouchel, and J. Xia (2021): "The Future of Labor: Automation and the Labor Share in the Second Machine Age," Discussion paper, Mimeo, FRB Philadelphia.

Chernoff, A. W., and C. Warman (2020): "COVID-19 and Implications for Automation," Working Paper 27249, National Bureau of Economic Research.

Cho, J., T. DeStefano, H. Kim, and J. Paik (2021): "Closing the Italian digital gap: The role of skills, intangibles and policies," Discussion paper, Mimeo, Harvard Business School.

Chow, M. C., T. C. Fort, C. Goetz, N. Goldschlag, J. Lawrence, E. R. Perlman, M. Stinson, and T. K. White (2021): "Redesigning the Longitudinal Business Database," Discussion paper, National Bureau of Economic Research.

Comin, D. A., M. Cruz, X. Cirera, K. M. Lee, and J. Torres (2022): "Technology and Resilience," Working Paper 29644, National Bureau of Economic Research.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2021): "The adjustment of labor markets to robots," *Journal of the European Economic Association*, 19(6), 3104–3153.

de Souza, G., and A. Sollaci (2020): "Robots, Tools, and Jobs: Evidence from Brazilian Labor Markets,"

Discussion paper, Mimeo, The University of Chicago.
Dechezlepretre, A., D. Hemous, M. Olsen, and C. Zanella (2021): "Induced Automation: Evidence from Firm-level Patent Data," Discussion paper, Mimeo, Zurich University.
Dechezleprêtre, A., D. Hémous, M. Olsen, and C. Zanella (2021): "Induced automation: evidence from firm-level patent data," Discussion Paper 384, University of Zurich, Department of Economics, Working Paper.
Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2020): "Changing Business Dynamism and Productivity: Shocks versus Responsiveness," *American Economic Review*, (forthcoming).
Dinlersoz, E. M., and Z. Wolf (2018): "Automation, Labor Share, and Productivity: Plant-Level Evidence from U.S. Manufacturing," Discussion paper, Center for Economic Studies Working Paper CES-WP-18-39.
Dixon, J., B. Hong, and L. Wu (2020): "The employment consequences of robots: Firm-level evidence," Discussion paper, Mimeo, Statistics Canada.
Doms, M., T. Dunne, and K. Troske (1997): "Workers, Wages, and Technology," *Quarterly Journal of Economics*, 112(1), 253–290.
Dunne, T., L. Foster, J. Haltiwanger, and K. R. Troske (2004): "Wage and productivity dispersion in United States manufacturing: The role of computer investment," *Journal of Labor Economics*, 22(2), 397–429.
Eckert, F., S. Ganapati, and C. Walsh (2020): "Skilled Scalable Services: The New Urban Bias in Economic Growth," Discussion paper, CESifo Working Paper No. 8705.
Ford, M. (2015): *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books, New York.
Fort, T. C., J. R. Pierce, and P. K. Schott (2018): "New Perspectives on the Decline of US Manufacturing Employment," *Journal of Economic Perspectives*, 32(2), 47–72.
Foster, L., J. Haltiwanger, S. Klimek, C. Krizan, and S. Ohlmacher (2016): "The Evolution of National Retail Chains: How We Got Here," in *Handbook on the Economics of Retailing and Distribution*, ed. by E. Basker. Elgar.
Foster, L., J. C. Haltiwanger, and C. J. Krizan (2001): *Aggregate Productivity Growth: Lessons from Microeconomic Evidence*pp. 303–372. University of Chicago Press.
Genz, S., T. Gregory, M. Janser, F. Lehmer, and B. Matthes (2021): "How Do Workers Adjust When Firms Adopt New Technologies?," Discussion paper, IZA DP No. 14626.
Graetz, G., and G. Michaels (2018): "Robots at Work," *The Review of Economics and Statistics*, 100(5), 753–768.
Hubmer, J., and P. Restrepo (2021): "Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital-Labor Substitution," Working Paper 28579, National Bureau of Economic Research.
Humlum, A. (2020): "Robot Adoption and Labor Market Dynamics," Working paper, University of Chicago.
Kehrig, M., and N. Vincent (2021): "The Micro-Level Anatomy of the Labor Share Decline*," *The Quarterly Journal of Economics*, 136(2), 1031–1087.
Koch, M., I. Manuylov, and M. Smolka (2021): "Robots and firms," *The Economic Journal*, 131(638), 2553–2584.
Kogan, L., D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller (2021): "Technology-Skill Complementarity and Labor Displacement: Evidence from Linking Two Centuries of Patents with Occupations," Working Paper 29552, National Bureau of Economic Research.
Leigh, N. G., and B. R. Kraft (2018): "Emerging robotic regions in the United States: insights for regional economic evolution," *Regional Studies*, 52(6), 804–815.
Loucks, J., T. Davenport, and D. Schatsky (2018): *State of AI in the Enterprise, 2nd Edition*. Deloitte Insights.
Mann, K., and L. Püttmann (2019): "Benign effects of automation: New evidence from patent texts," *The Review of Economics and Statistics*, pp. 1–45.
Martinez, J., and J. Moen-Vorum (2021): "Automation Potential and Diffusion," Discussion Paper 384, Mimeo, London Business School.
McElheran, K. (2015): "Do market leaders lead in business process innovation? The case (s) of e-business

adoption," *Management science*, 61(6), 1197–1216.

McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, and N. Zolas (2022): "AI Adoption in America: Who, What and Where," Working paper, US Census Bureau, Center for Economic Studies.

Perla, J., C. Tonetti, and M. E. Waugh (2021): "Equilibrium Technology Diffusion, Trade, and Growth," *American Economic Review*, 111(1), 73–128.

Rodrigo, R. (2021): "Robot Adoption, Organizational Capital, and the Productivity Paradox," Discussion paper, Mimeo, Georgetown University.

Schwab, K. (2017): *The fourth industrial revolution*. Currency.

Susskind, R. E., and D. Susskind (2015): *The future of the professions: How technology will transform the work of human experts*. Oxford University Press, USA.

Zolas, N., Z. Kroff, E. Brynjolfsson, K. McElheran, D. Beede, C. Buffington, N. Goldschlag, L. Foster, and E. Dinlersoz (2020a): "Measuring Technology Adoption in Enterprise-Level Surveys: The Annual Business Survey," Discussion paper, US Census.

Zolas, N., Z. Kroff, E. Brynjolfsson, K. McElheran, D. N. Beede, C. Buffington, N. Goldschlag, L. Foster, and E. Dinlersoz (2020b): "Advanced Technologies Adoption and Use by US Firms: Evidence from the Annual Business Survey," Discussion paper, National Bureau of Economic Research.

Appendix A: Development of the ABS 2019 Module

Work on the development of the 2019 ABS module began in the Spring of 2018. The questions for the module were developed over a period of several months in cooperation with NCSES, and with input from economists at Massachusetts Institute of Technology and Boston University. The initial technologies included in the 2019 ABS module consisted of those relevant for automation: specialized software, specialized equipment, robotics, and artificial intelligence. Cloud computing was added later in the process, as it complements some of these technologies, particularly AI, and facilitates automation. The initial draft of the module only considered technology adoption within the context of the processes for producing goods or services. When confronted with the fact that there is little up-to-date information on which firms actually provide these technologies as their products and services, it was decided to duplicate the questions for firms which identify themselves as sellers of the goods or services embedding the technologies (e.g. providers of machine learning software, or robot producers).

Asking about the exact adoption date of the technologies, a key piece of information for pinning down the timing of adoption in the life-cycle of a firm, would demand a high degree of recall from respondents. Instead, the survey measures adoption and use during a specified window of time, 2016–2018—the reference period. Similarly, questions on the exact magnitude of change in worker counts attributable to technology use impose high respondent burden. It was decided that inquiring about only the direction of change in worker counts was a more viable option. Cognitive testing of the module on a sample of potential respondents took place in late summer and fall of 2018. The testing process revealed some minor changes to the definitions of each of the technologies to make them more transparent for the respondents, and streamlined parts of the module. In the end, the module settled on five questions, all of which are asked for each technology separately.

Appendix B: Additional empirical results for the 2019 ABS

This appendix provides additional empirical results discussed in the main text. The following table of contents summarizes the material included in this appendix, and links it to the relevant sections in the main text.

- Tables A-1 and A-2 reproduce the findings from Tables 4 and 5 separately for manufacturing and non-manufacturing firms. Table A-3 reproduces Tables 4 and 5 but now uses as outcome a dummy for whether firms adopted the technology for automation. These tables explain firm adoption rates as a function of size and age classes, detailed industry dummies, employment shares by state, and labor productivity and labor shares from the 2015 LBD.

- Table A-4 reproduces the results from Table 6 in the main text, but now explains labor productivity as a function of size classes defined by lagged employment.

- Figure A-1 reports firms assessments on the effects of technology on their employment level and demand for skills by size. The estimates come from a generalized ordered logit model (controlling for size, age, and sector).

- Figure A-2 reports firms assessments on the effects of technology on their employment level and demand for skills by age. The estimates come from a generalized ordered logit model (controlling for size, age, and sector).

Table A-1: Regressions explaining the adoption of advanced technologies in manufacturing, ABS data for 2016–2018.

Dependent variable:	Artificial intelligence		Robotics		Dedicated equipment		Specialized software		Cloud computing	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Employment percentile 50th–90th	0.016	0.015	0.064	0.065	0.166	0.159	0.185	0.176	0.153	0.145
	(0.003)	(0.003)	(0.004)	(0.004)	(0.007)	(0.007)	(0.008)	(0.008)	(0.007)	(0.007)
Employment percentile 90th–95th	0.027	0.026	0.145	0.146	0.269	0.261	0.284	0.274	0.212	0.204
	(0.006)	(0.006)	(0.011)	(0.011)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Employment percentile 95th–99th	0.046	0.046	0.238	0.239	0.340	0.334	0.340	0.332	0.297	0.290
	(0.007)	(0.007)	(0.012)	(0.012)	(0.016)	(0.016)	(0.015)	(0.015)	(0.016)	(0.016)
Employment percentile 99th+	0.107	0.107	0.396	0.398	0.408	0.404	0.370	0.366	0.378	0.375
	(0.015)	(0.015)	(0.023)	(0.023)	(0.023)	(0.023)	(0.024)	(0.024)	(0.025)	(0.025)
Age percentile 10th–50th	-0.009	-0.009	-0.016	-0.016	-0.043	-0.043	-0.041	-0.041	-0.060	-0.060
	(0.005)	(0.005)	(0.006)	(0.006)	(0.012)	(0.012)	(0.012)	(0.012)	(0.011)	(0.011)
Age percentile 50th–75th	-0.012	-0.012	-0.027	-0.027	-0.068	-0.070	-0.058	-0.061	-0.099	-0.102
	(0.005)	(0.005)	(0.007)	(0.007)	(0.012)	(0.012)	(0.013)	(0.013)	(0.012)	(0.012)
Age percentile 75th–95th	-0.011	-0.011	-0.027	-0.027	-0.088	-0.092	-0.096	-0.101	-0.123	-0.127
	(0.005)	(0.005)	(0.008)	(0.008)	(0.013)	(0.013)	(0.015)	(0.015)	(0.014)	(0.014)
Age percentile 95th+	-0.003	-0.003	0.001	0.001	-0.022	-0.023	0.018	0.017	0.014	0.012
	(0.005)	(0.005)	(0.008)	(0.008)	(0.012)	(0.012)	(0.014)	(0.014)	(0.013)	(0.013)
Log of labor productivity	0.006		0.021		0.034		0.044		0.039	
	(0.002)		(0.002)		(0.004)		(0.005)		(0.005)	
Log of labor share		-0.004		-0.022		-0.019		-0.025		-0.023
		(0.003)		(0.003)		(0.005)		(0.007)		(0.006)
Log of average wage		0.008		0.020		0.055		0.072		0.062
		(0.002)		(0.003)		(0.006)		(0.006)		(0.006)
R-squared	1.7%	1.7%	10%	9.9%	8.5%	8.6%	9.8%	10%	8.9%	9.1%
Observations	22,500	22,500	23,500	23,500	23,000	23,000	22,500	22,500	23,000	23,000

Notes: The table reports results from a regression of firm-level adoption on firm characteristics, 6-digit industry dummies, and employment shares by state. The sample includes firms in the ABS in the manufacturing sector. Columns 1, 2 report results for the adoption of artificial intelligence. Columns 3, 4 report results for the adoption of robotics. Columns 5, 6 report results for the adoption of dedicated equipment. Columns 7, 8 report results for the adoption of specialized software. Columns 9, 10 report results for the adoption of cloud computing. To protect confidentiality, these table uses a coarser definition of size and age brackets than Tables 4 and 5.

Table A-2: Regressions explaining the adoption of advanced technologies outside of manufacturing, ABS data for 2016–2018.

Dependent variable:	Artificial intelligence		Robotics		Dedicated equipment		Specialized software		Cloud computing	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Employment percentile 50th–90th	0.010	0.009	0.008	0.008	0.068	0.067	0.143	0.137	0.130	0.124
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
Employment percentile 90th–95th	0.016	0.016	0.020	0.020	0.105	0.102	0.236	0.226	0.225	0.214
	(0.003)	(0.003)	(0.003)	(0.003)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)
Employment percentile 95th–99th	0.015	0.014	0.019	0.019	0.116	0.113	0.265	0.254	0.272	0.260
	(0.004)	(0.004)	(0.003)	(0.003)	(0.009)	(0.009)	(0.011)	(0.011)	(0.010)	(0.010)
Employment percentile 99th+	0.029	0.029	0.031	0.031	0.163	0.163	0.314	0.310	0.330	0.326
	(0.005)	(0.005)	(0.004)	(0.004)	(0.022)	(0.022)	(0.022)	(0.023)	(0.022)	(0.022)
Age percentile 10th–50th	-0.001	-0.001	-0.002	-0.002	-0.007	-0.007	-0.013	-0.013	-0.026	-0.027
	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)
Age percentile 50th–75th	-0.005	-0.005	-0.004	-0.004	-0.026	-0.026	-0.039	-0.040	-0.068	-0.070
	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)
Age percentile 75th–95th	-0.007	-0.007	-0.007	-0.007	-0.028	-0.029	-0.040	-0.043	-0.085	-0.088
	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)
Age percentile 95th+	-0.005	-0.006	-0.006	-0.006	-0.025	-0.026	-0.048	-0.051	-0.073	-0.076
	(0.003)	(0.003)	(0.002)	(0.002)	(0.008)	(0.008)	(0.009)	(0.009)	(0.008)	(0.008)
Log of labor productivity	0.002		0.002		0.012		0.040		0.043	
	(0.001)		(0.001)		(0.002)		(0.002)		(0.002)	
Log of labor share		-0.001		-0.001		-0.009		-0.028		-0.030
		(0.001)		(0.001)		(0.002)		(0.002)		(0.002)
Log of average wage		0.003		0.002		0.016		0.056		0.060
		(0.001)		(0.001)		(0.002)		(0.003)		(0.003)
R-squared	1.7%	1.7%	3.6%	3.6%	11.1%	11.1%	14.1%	14.3%	14.3%	14.4%
Observations	94,500	94,500	97,000	97,000	95,000	95,000	94,000	94,000	95,000	95,000

Notes: The table reports results from a regression of firm-level adoption on firm characteristics, 6-digit industry dummies, and employment shares by state. The sample includes firms in the ABS outside of the manufacturing sector. Columns 1, 2 report results for the adoption of artificial intelligence. Columns 3, 4 report results for the adoption of robotics. Columns 5, 6 report results for the adoption of dedicated equipment. Columns 7, 8 report results for the adoption of specialized software. Columns 9, 10 report results for the adoption of cloud computing. To protect confidentiality, these table uses a coarser definition of size and age brackets than Tables 4 and 5.

Table A-3: Regressions explaining the adoption of advanced technologies for automation, ABS data for 2016–2018.

Dependent variable:	Artificial intelligence		Robotics		Dedicated equipment		Specialized software		Cloud computing	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Employment percentile 50th–75th	0.003	0.003	0.003	0.004	0.022	0.022	0.060	0.058	0.041	0.040
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.004)	(0.004)	(0.003)	(0.003)
Employment percentile 75th–90th	0.006	0.006	0.008	0.009	0.040	0.040	0.099	0.097	0.073	0.071
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.005)	(0.005)	(0.004)	(0.004)
Employment percentile 90th–95th	0.009	0.009	0.016	0.016	0.055	0.055	0.146	0.142	0.107	0.104
	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)	(0.009)	(0.009)	(0.007)	(0.007)
Employment percentile 95th–99th	0.011	0.012	0.018	0.018	0.052	0.052	0.156	0.152	0.111	0.108
	(0.003)	(0.003)	(0.002)	(0.002)	(0.005)	(0.005)	(0.011)	(0.011)	(0.008)	(0.008)
Employment percentile 99th+	0.022	0.022	0.033	0.033	0.079	0.080	0.235	0.234	0.147	0.146
	(0.004)	(0.004)	(0.004)	(0.004)	(0.010)	(0.010)	(0.025)	(0.025)	(0.015)	(0.015)
Age percentile 10th–50th	-0.003	-0.003	-0.000	-0.000	-0.008	-0.008	-0.022	-0.022	-0.020	-0.020
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.005)	(0.005)	(0.004)	(0.004)
Age percentile 50th–75th	-0.005	-0.005	-0.000	-0.000	-0.015	-0.015	-0.035	-0.036	-0.040	-0.041
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.005)	(0.005)	(0.004)	(0.004)
Age percentile 75th–90th	-0.005	-0.005	-0.001	-0.001	-0.015	-0.015	-0.041	-0.042	-0.050	-0.051
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.006)	(0.006)	(0.005)	(0.005)
Age percentile 90th–95th	-0.006	-0.005	.		-0.022	-0.022	-0.037	-0.038	-0.062	-0.063
	(0.002)	(0.002)			(0.005)	(0.005)	(0.009)	(0.009)	(0.006)	(0.006)
Age percentile 95th–99th	-0.009	-0.008	0.000	0.001	-0.031	-0.031	-0.075	-0.076	-0.080	-0.081
	(0.002)	(0.002)	(0.001)	(0.001)	(0.005)	(0.005)	(0.011)	(0.011)	(0.010)	(0.010)
Age percentile 95th+	-0.006	-0.005	.		-0.016	-0.016	-0.045	-0.046	-0.059	-0.060
	(0.002)	(0.002)			(0.005)	(0.005)	(0.009)	(0.009)	(0.006)	(0.006)
Log of labor productivity	0.001		0.002		0.005		0.017		0.015	
	(0.000)		(0.000)		(0.001)		(0.002)		(0.001)	
Log of labor share		-0.001		-0.002		-0.005		-0.013		-0.012
		(0.000)		(0.000)		(0.001)		(0.002)		(0.002)
Log of average wage		0.001		0.001		0.005		0.023		0.020
		(0.000)		(0.000)		(0.001)		(0.002)		(0.002)
R-squared	0.8%	0.8%	4.8%	4.8%	5.4%	5.4%	10%	10.1%	7.4%	7.5%
Observations	114,000	114,000	118,000	118,000	96,500	96,500	74,000	74,000	79,000	79,000

Notes: The table reports results from a regression of firm-level adoption of advanced technologies for automation on firm characteristics, 6-digit industry dummies, and employment shares by state. The sample includes firms that adopted technologies for automation and firms that did not adopt the technology. Columns 1, 2 report results for the adoption of artificial intelligence. Columns 3, 4 report results for the adoption of robotics. Columns 5, 6 report results for the adoption of dedicated equipment. Columns 7, 8 report results for the adoption of specialized software. Columns 9, 10 report results for the adoption of cloud computing. To protect confidentiality, in columns 3–4 the coefficients for firms in the age percentiles 75th to 90th and 90th to 95th are pooled together, as well as firms in the age percentiles 95th to 99th and above the 99th percentile. These coefficients are reported under the row corresponding to firms in the 75th to 90th and 95th to 99th age percentiles, respectively.

Table A-4: Regressions explaining firm labor productivity in 2019 as a function of technology adoption and lagged firm size, ABS data for 2016–2018.

Sector:	Dependent Variable: Log of Labor Productivity in 2019 All sectors combined			
	(1)	(2)	(3)	(4)
Lagged employment percentile 0th–50th	0.099	0.116	0.113	0.115
	(0.008)	(0.008)	(0.008)	(0.008)
Lagged employment percentile 75th–90th	0.018	0.005	0.009	0.007
	(0.010)	(0.009)	(0.009)	(0.009)
Lagged employment percentile 90th–95th	0.068	0.045	0.051	0.049
	(0.015)	(0.015)	(0.015)	(0.015)
Lagged employment percentile 95th–99th	0.111	0.083	0.092	0.087
	(0.018)	(0.018)	(0.018)	(0.018)
Lagged employment percentile 99th+	0.264	0.225	0.238	0.232
	(0.041)	(0.040)	(0.040)	(0.040)
Age percentile 10th–50th	0.002	0.008	0.006	0.007
	(0.016)	(0.016)	(0.016)	(0.016)
Age percentile 50th–75th	-0.032	-0.019	-0.022	-0.021
	(0.016)	(0.016)	(0.016)	(0.016)
Age percentile 75th–90th	-0.055	-0.039	-0.043	-0.041
	(0.017)	(0.017)	(0.017)	(0.017)
Age percentile 90th–95th	-0.090	-0.072	-0.077	-0.075
	(0.022)	(0.022)	(0.022)	(0.022)
Age percentile 95th–99th	-0.144	-0.117	-0.126	-0.121
	(0.033)	(0.032)	(0.032)	(0.032)
Age percentile 99th+	-0.051	-0.035	-0.039	-0.037
	(0.021)	(0.021)	(0.021)	(0.021)
Artificial intelligence		-0.001		
		(0.019)		
Cloud computing		0.101		
		(0.009)		
Robotics		0.091		
		(0.021)		
Specialized software		0.076		
		(0.009)		
Dedicated equipment		-0.022		
		(0.010)		
Technology user			0.128	
			(0.007)	
One technology				0.083
				(0.011)
Two technologies				0.143
				(0.009)
Three technologies				0.152
				(0.011)
Four technologies				0.193
				(0.021)
Five technologies				0.238
				(0.043)
R-squared	30.4%	30.9%	30.7%	30.8%
Observations	101,000	101,000	101,000	101,000

Notes: The table reports results from a regression of firm labor productivity on size and age groups, 6-digit industry dummies, employment shares by state, and measures of technology adoption from the 2019 ABS. Size groups are based on lagged employment percentiles. To protect confidentiality, the number of observations has been rounded.

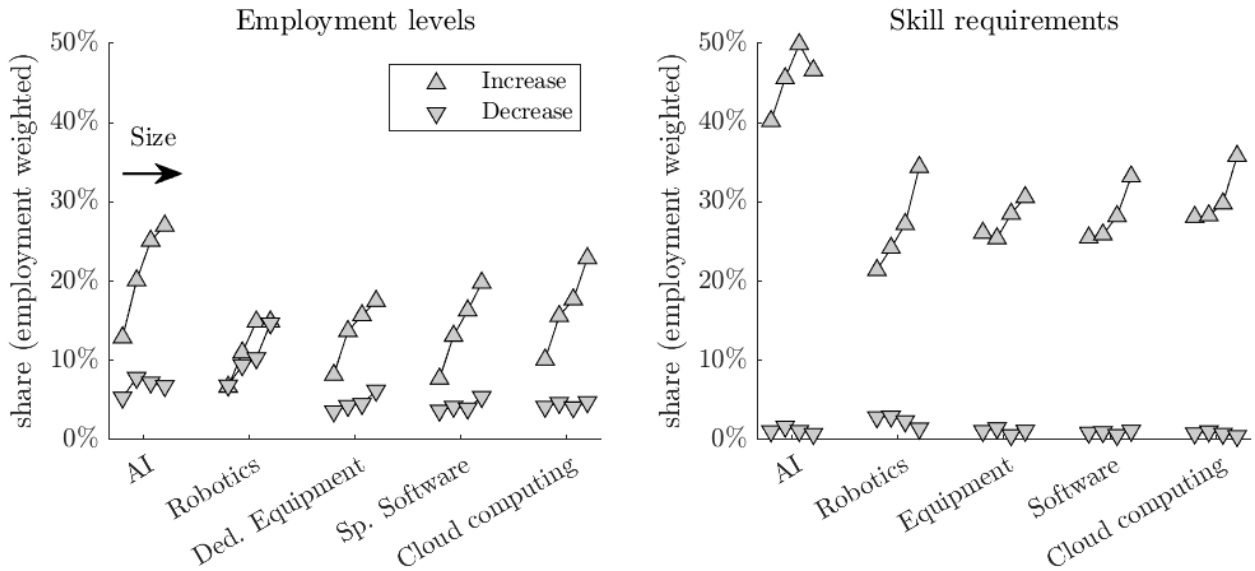

Figure A-1: Reported changes in employment levels and skill demand by firms adopting advanced technologies, employment-weighted shares by size from 2019 ABS. The markers provide the employment-weighted responses for firms with 0–9 workers, 10–49 workers, 50–249 workers, and more than 250 workers. The estimates come from a generalized ordered logit model (controlling for size, age, and sector).

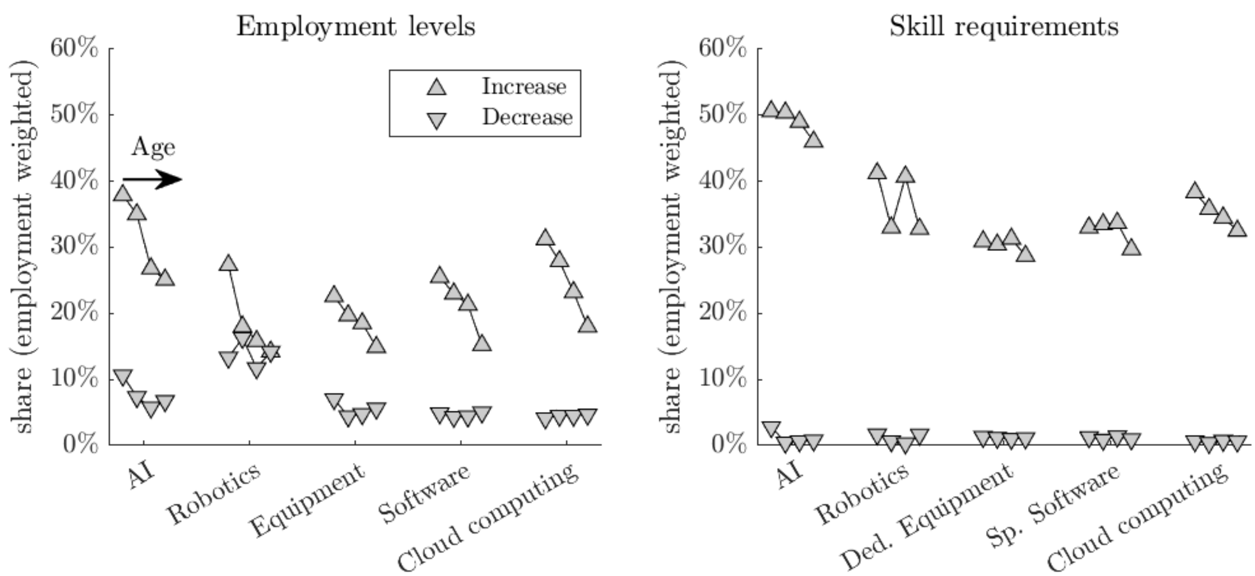

Figure A-2: Reported changes in employment levels and skill demand by firms adopting advanced technologies, employment-weighted shares by age from 2019 ABS. The markers provide the employment-weighted responses for firms of 0–5 years, 6–10 years, 11–20 years, and more than 21 years. The estimates come from a generalized ordered logit model (controlling for size, age, and sector).

Appendix C: Life-cycles of Adopters and Selection

The previous sections documented sizable differences in size (measured in employment) between adopters and non adopters. To further explore the source of these differences in size, we now study the employment history of adopters and compare it to that of non-adopters in their same industry and cohort. To do so, we create a pseudo-firm establishment panel using the LBD by matching all establishments in the LBD whose firm identifiers match those in the ABS. The working assumption is that tracking these establishments back in time provides information on the life-cycle dynamics of their contemporary parent firm in 2016–2018.[37] Our approach misses the fact that some establishments that were part of a firm in the past might have changed parent firms over time due to mergers and acquisitions or that some establishments that were part of a firm might have exited by 2016. Our approach is also only able to study the life cycle of establishments that are part of adopters and non-adopters that survived until 2016–2018, which might introduce non-trivial selection issues.[38]

To illustrate the structure of the pseudo-firm establishment panel, we first focus on the case of robotics and plot average establishment size by cohort and age separately for adopters and non-adopters.[39] The results of this exercise are shown in Figure A-3, which provides the life cycle employment profile for each cohort of establishments in adopting firms and non-adopting firms. As expected, we can only trace a smaller part of the life cycle of more recent cohorts.[40]

[37]One alternative approach involves using firm identifiers from the LBD to match a changing set of establishments over time to each firm in the ABS. However, firm identifiers in the LBD cannot be treated as longitudinally consistent in the same way establishment-level identifiers are (Chow et al., 2021). See also Foster et al. (2016) for distinct but related tracking of activity of firms backwards in time.

[38]For example, non-adopters from early cohorts that make it to 2018 are more positively selected than non adopters from more recent cohorts, and this selection might affect adopters and non-adopters differently.

[39]For our measure of size we use the inverse hyperbolic sine (ihs) transformation. The inverse hyperbolic sine approximates the log transformation but permits the inclusion of zeros. $ihs(emp) = ln(emp + (1 + emp^2)^{1/2}$. Though by conditioning on survival to 2018 we expect far fewer zero-employment observations, establishments do, however infrequently in our sample, transition in and out of positive employment–deactivating and reactivating over time.

[40]The LBD includes a series of establishments that were observed for the first time in 1976—the first year in the LBD—, but whose exact age is unknown. This "left-censored" establishments are pooled together and not included in the figure. Among the left-censored establishments, we also find a that those belonging to adopter firms are 60% larger than those in non adopting firms. See Chow et al. (2021) for details on the treatment of firm and establishment age in the LBD.

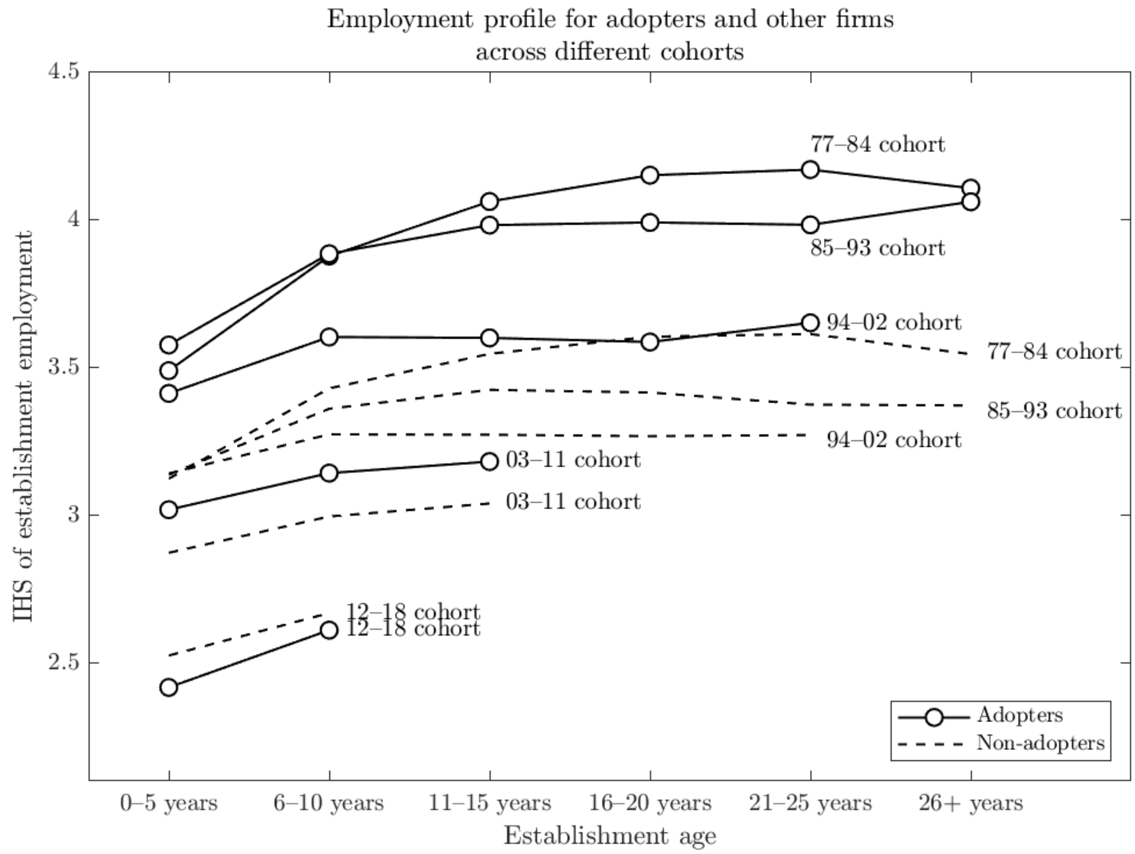

Figure A-3: Life-cycle employment profile for robot adopters and non-adopters of different cohorts, based on firm adoption status from the 2019 ABS.

The figure reveals three key patterns. First, for all but the most recent cohort, establishments that belong to adopting firms are larger at all ages than establishments in non-adopting firms. Second, these differences were already manifest at an early age and become more pronounced as firms age. That is, establishments that belong to adopting firms tend to start larger and have a steeper size-age life-cycle profile than establishments from their same cohort that belong to non-adopters. Finally, the size gap between establishments at adopting and non-adopting firms has been falling over time.

To explore whether these patters hold for other technologies and how significant they are, we now turn to a second regression analysis, where we explain establishment employment as a function of cohort and age effects that might vary differentially for adopters. In particular, we estimate the regression model

$$\text{IHS of employment}_{e,f,i,c,a,t} = \alpha_a + \alpha_a^A \times \text{Adopter}_f^{\text{tech}} + \beta_c + \beta_c^A \times \text{Adopter}_f^{\text{tech}} + \gamma_{i,t} + \epsilon_{e,f,i,c,a,t}, \quad \text{(A-4)}$$

for an establishment e in parent firm f (assigned based on the 2019 LBD), industry i, cohort c, age a, and in calendar year t. This model explains establishment employment as a function of age dummies α_a, cohort dummies β_c^A, and industry by calendar year dummies $\gamma_{i,t}$, which accounts for differences in macroeconomic trends across 4-digit industries.[41] In addition, we allow establishments from adopting firms to have different life-cycle employment profiles $\alpha_a^A \times \text{Adopter}_f^{\text{tech}}$ and cohort effects $\beta_c^A \times \text{Adopter}_f^{\text{tech}}$. The dummies α_a^A capture differences in establishments' life cycle profiles between adopters and non-adopters. The dummies β_c^A capture the difference in initial size between establishments at adopting and non-adopting firms for each cohort.[42]

Figure A-4 provides the estimates for the differences in initial size by cohort β_c^d (shown in Panel A) and differences in employment by age (shown in Panel B) between establishments at adopting and non-adopting firms. In line with our inspection of Figure A-3, we find large and statistically significant differences in initial employment between adopters and non-adopters' establishments. These size differences have also been falling over time. For example, establishments at firms using robotics were 30% larger than others in the 1977–1984 cohort, relative to the size difference observed for the 2012–2018 cohort (the excluded category). On the other hand, establishments at firms using robotics were 8% larger than others in the 2003–2011 cohort, relative to the size difference observed for the 2012–2018 cohort (the excluded category). Finally, and with the exception of artificial intelligence, we also find that establishments at adopting firms are on a more steep life-cycle growth trajectory, though these differences are smaller than those explained by their initial size and captured by the differential cohort effects. For example, establishments in firms using robotics grow 10% more over the first 20 years of their lives than establishments at non-adopters.[43]

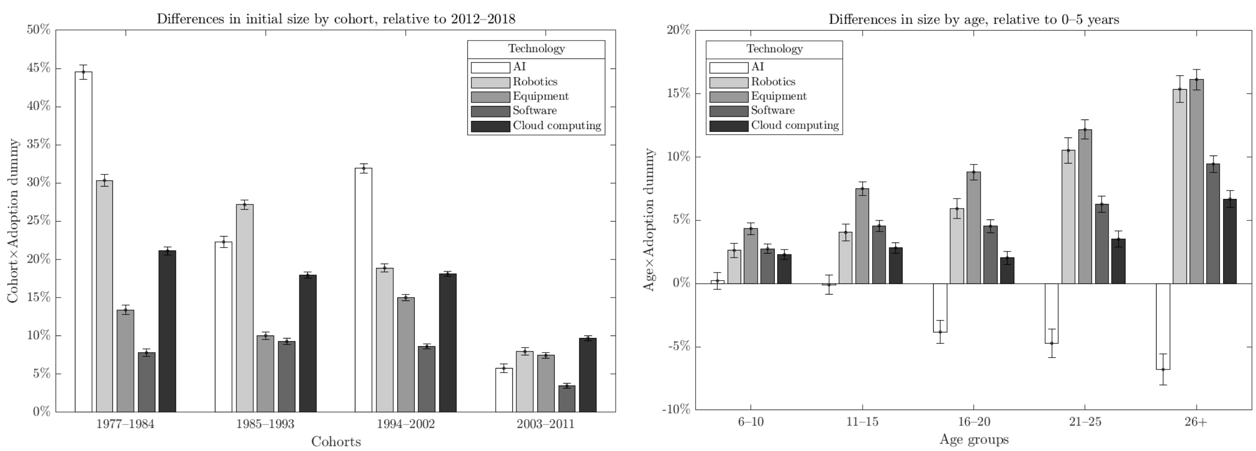

Panel A: Cohort×Adopt Effects

Panel B: Age×Adopt Effects

Figure A-4: Estimates of differences in initial size by cohort (Panel A) and size by age (Panel B) between adopters and non adopters.

One important challenge when interpreting these facts is that we do not know the exact date

[41] Industry codes are allowed to vary within an establishment over time in the relatively rare case an establishment's industry changes.

[42] We cannot allow for differential effects by age, cohort and time, since these variables are collinear.

[43] Though not reported in Figure A-4, we also saturate the model with dummies for the left-censored establishments. We find that left-censored establishments in adopting firms are 20–40% larger than left-censored establishments in other firms once we control for industry dummies. These differences might reflect a combination of differences in their initial sizes and the steeper life-cycle profile for adopters.

when adoption took place, and so it is hard to infer whether differences in size or the more steep life-cycle profile of adopters is caused by technology, or, on the other hand, if these were the factors that favored the adoption of advanced technology in the first place. However, one plausible interpretation that aligns with firms' assessment of the effects of these technologies on employment levels and the high reported cost of adopting these technologies is that these patterns reflect the selection of large and fast-growing firms into using advanced technologies for automation. In fact, most of the size differences depicted in Figure A-3 preceded the diffusion of industrial robots and other automation technologies among US firms. The view that emerges is that adoption entails high fixed costs, which cause firms that were initially larger or that were on a more favorable growth trajectory to select into adopting these technologies. Over time, as advanced technologies diffuse and become standardized and mature, the cost of adoption might have declined, explaining why size differences between adopters and non adopters are shrinking for more recent cohorts.